RM Methods for Multiple Fare Structure Environments

By

Matthew R. Kayser

B.S., Operations Research
United States Air Force Academy, 2006

SUBMITTED TO THE DEPARTMENT OF CIVIL AND ENVIRONMENTAL
ENGINEERING IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE
DEGREE OF

MASTER OF SCIENCE IN TRANSPORTATION
AT THE
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

JUNE 2008

Signature of Author: ……...……………...…………………………………………..…….
Department of Civil and Environmental Engineering
May 9, 2008

Certified by: …………………………..…….………………………………………….
Peter P. Belobaba
Principal Research Scientist of Aeronautics and Astronautics
Thesis Supervisor

Certified by: …………….…………...…………………………………………………..
Cynthia Barnhart
Associate Dean of Engineering for Academic Affairs
Professor of Civil and Environmental Engineering
And Co-Director, Operations Research Center
Thesis Reader

Accepted by: ………………….…………...……………………………………………
Daniele Veneziano
Chairman, Departmental Committee for Graduate Students

1. REPORT DATE **01 JUN 2008**	2. REPORT TYPE **N/A**	3. DATES COVERED **-**	
4. TITLE AND SUBTITLE **RM Methods for Multiple Fare Structure Environments**		5a. CONTRACT NUMBER	
		5b. GRANT NUMBER	
		5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)		5d. PROJECT NUMBER	
		5e. TASK NUMBER	
		5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Massachusetts Institute Of Technology**		8. PERFORMING ORGANIZATION REPORT NUMBER **CI08-0018**	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) **The Department of the Air Force AFIT/ENEL, Bldg 16 2275 D Street WPAFB, OH 45433**		10. SPONSOR/MONITOR'S ACRONYM(S)	
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT **Approved for public release, distribution unlimited**			
13. SUPPLEMENTARY NOTES **The original document contains color images.**			
14. ABSTRACT			
15. SUBJECT TERMS			

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT **UU**	18. NUMBER OF PAGES **116**	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**			

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std Z39-18

RM Methods for Multiple Fare Structure Environments

By

Matthew R. Kayser

Submitted to the Department of Civil and Environmental Engineering on May 9, 2008
in Partial Fulfillment of the Requirements for the Degree of Master of Science in Transportation

ABSTRACT

The rapid growth of Low Cost Carriers (LCC) and their simplified fare structures has created "semi-restricted" fare structures where lower classes are undifferentiated except for price, while higher fare classes are still differentiated by various advance purchase and booking restrictions. The problem this causes is two-fold: first, traditional revenue management systems, which operate based on the assumption of demand independence, will see demand "spiral down" into the lowest booking classes as passengers will buy the lowest available fare in the absence of fare restrictions. Second, airlines must maximize network revenues across two different fare structures, a more-restricted structure used on markets without an LCC presence, and the semi-restricted structure for markets where LCC competition exists.

This thesis describes methods of dealing with these two problems: Hybrid Forecasting (HF), which forecasts "product-oriented" demand using traditional forecasting methods while simultaneously forecasting "price-oriented" demand for those passengers who will buy the lowest available fare, and Fare Adjustment (FA), which is used at the booking limit optimizer level to account for the sell-up potential of passengers (probability a passenger will book in a higher class if his original booking class is denied). Fare Adjustment allows the airline to deal with multiple fare structures separately. The goal of this thesis is to provide a comprehensive summary of results when an airline uses HF and FA simultaneously in two different multiple fare structure, competitive networks. An alternate Fare Adjustment formulation will also be introduced and tested in these competitive environments.

Results from the Passenger Origin-Destination Simulator (PODS) demonstrate that in a more restrictive network, HF and FA used in conjunction with one another achieve revenue increases of approximately 2-4% above traditional forecasting methods. In an environment with a fully unrestricted fare structure for LCC markets, HF and FA together generate revenue gains of over 20% above traditional methods.

Thesis Supervisor: Peter P. Belobaba
Title: Principal Research Scientist of Aeronautics and Astronautics

Thesis Reader: Cynthia Barnhart
Title: Associate Dean of Engineering for Academic Affairs
     Professor of Civil and Environmental Engineering
     And Co-Director, Operations Research Center

<u>Acknowledgements</u>

TABLE OF CONTENTS

LIST OF FIGURES

5.2   NETWORK S4

6   CONCLUSIONS

List of Tables

CHAPTER 1
INTRODUCTION

Although it has confused business and leisure travelers alike, Revenue Management (RM) has played a large role in the profitability and sustainability of airlines worldwide. Perhaps the best and most succinct definition of RM comes from American Airlines, who in their 1987 annual report described RM as "selling the right seats to the right customers at the right prices"[1]. Even after 20 years, airlines are still working toward this seemingly simplistic goal in order to maximize revenues.

Weatherford[2] more accurately describes the airline situation as Perishable-Asset Revenue Management (PARM). When a flight takes off with empty seats, the revenue opportunity for those seats is lost forever, thus leading to the assertion that airlines oversee the most perishable inventory in the world. The task facing an airline is how best to allocate these perishable seats on an individual flight leg to passengers flying on multiple Origin-Destination (O-D) itineraries, possibly from different fare structures, in order to maximize total network revenue. The ultimate (albeit unattainable) goal of RM is to sell passengers seats at their maximum willingness-to-pay (WTP).

The rest of this chapter will provide a brief history of revenue management in the airline industry, as well as develop the need for and introduce RM techniques to be used to maximize revenues in the new fare structure environment encountered by the Network Legacy Carriers (NLC) due to the rise in prominence of Low Cost Carriers (LCC).

1.1   OVERVIEW OF REVENUE MANAGEMENT

The US airline industry was regulated by the Civil Aeronautics Board (CAB) until 1978, controlling both the routes carriers could fly and what fares they could charge on those routes[3]. However, airlines were able to perform a small portion of RM in the form of overbooking. They were able to forecast the number of passengers who would cancel reservations and simply not show up for the flight (*no-shows*), and therefore be able to calculate how many bookings above the actual capacity of the aircraft they would need to accept in order to achieve a high load factor without having large costs due to denied boardings (hotel and food vouchers, loss of customer goodwill, etc.).

---

[1] Smith, B.C., J.F. Leimkuhler, R.M. Darrow. (1992). Yield Management at American Airlines. *Interfaces.* Volume 22, Issue 1, pp. 8-31.
[2] Weatherford, L.R. (1991). Perishable Asset Revenue Management in General Business Situations, Ph.D. thesis, Darden Graduate School of Business Administration, University of Virginia, Charlottesville, VA.
[3] General Accounting Office. (1999). Airline Deregulation: Changes in Airfares, Service Quality, and Barriers to Entry. Report to Congressional Requesters. GAO/RCED-99-92. Washington, D.C.

The first commonly accepted example of airline RM came from overseas, in the United Kingdom. In the early 1970's BOAC (now known as British Airways) began offering two fares, with one being a lower fare for those passengers who were able to book twenty-one days in advance. This was the beginning of what was then called Yield Management (YM), as BOAC had to decide how many seats to keep available for the later booking passengers who would purchase the higher fare[4].

Back in the US, the CAB realized that the regulations on air travel were restricting the growth of the industry, and so in 1978 Congress deregulated the industry[3]. This sudden change brought about the rapid increase in RM activity, as airlines were now able to not only enter more markets, but were also free to charge based on demand in a market, not based on distance as was previously the case in the regulated environment[5]. Multiple airlines were able to offer service in the same markets, which gave passengers many more options as they were able to choose between multiple airlines and fare classes in a single market. As prices dropped due to this increase in competition, airlines were forced to develop different pricing structures to generate higher revenues by charging fares that were closer to an individual passenger's willingness-to-pay.

Ever since deregulation, airlines have focused on the two different types of passengers: business and leisure. Business travelers are characterized by purchasing close to departure, sensitivity to schedule and frequency, aversion to most fare restrictions (min stay, non-refundability, etc.), and willingness to pay higher prices as a result. On the other hand, leisure travelers normally plan ahead of time and are able to purchase tickets early in the booking process, are not sensitive to schedule or frequency, and are not as sensitive to fare restrictions, and are thus willing to buy lower priced tickets. In order for RM systems to be effective, they must have an accurate forecast of demand for these two unique sets of passengers.

Airlines found that in order to capture higher revenues and target business and leisure travelers separately, they would have to offer more than one fare class. Otherwise, they would have to operate under one of two conditions outlined in Figure 1-1: overprotection or dilution. If an airline's objective is to focus on high-yield business travelers and charge only a single high fare for the entire booking process, flights will leave with high yields but low load factors as there will be empty seats at this full-fare price. There will not be enough demand at this high price to fill seats and thus revenue will be lost due to the low number of bookings. The term "Yield Management" implies incorrectly that yield maximization is the goal, and has been gradually phased out in favor of "Revenue

---

[4] McGill, J.I., G.J. van Ryzin. (1999). Revenue Management: Research Overviews and Prospects. *Transportation Science*, Volume 33, Issue 2, pp. 233-256.
[5] Pickrell, D. (1991). The regulation and deregulation of US airlines. *Airline deregulation: international experiences.* Ed: Button, K. David Fulton Publishers, London, pp. 5-47.

Management", since a flight with a single full-fare passenger has the highest yield, but not the highest revenue[6].

Conversely, an airline can attempt to target low-paying leisure travelers by offering a single low fare throughout the booking process. This has the opposite effect, as flights will leave with high load factors but low yields. This is known as dilution, when revenue is lost because high-paying business travelers are allowed to buy fares far below their willingness-to-pay.
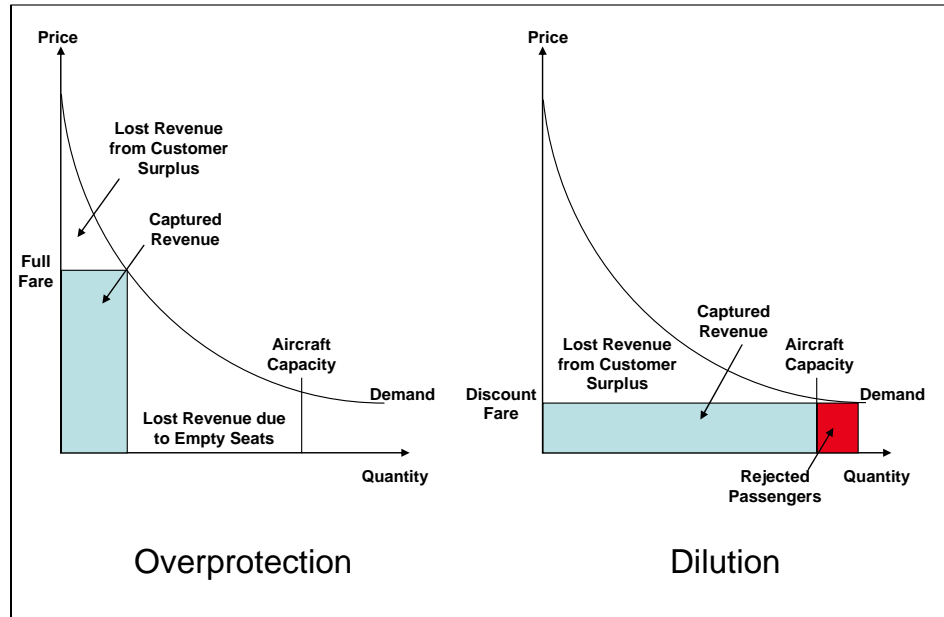
Figure 1-1: Revenue Losses due to Overprotection and Dilution

Obviously a single fare at any level does not maximize revenue, as there is untapped revenue either in the form of customer surplus or unused seats. Although it is unfeasible to charge each passenger his maximum WTP, airlines have found a way to group different sets of passengers together and charge each one of these groups a certain price, closer to the maximum WTP of each passenger in that group. By using business and leisure travelers' sensitivity to restrictions such as advanced purchase, Saturday night stay, cancellation fee, and non-refundability, airlines are able to segregate the two groups. This practice of offering differing products in terms of service and travel restrictions at different fare levels is called differential pricing[7]. Figure 1-2 demonstrates differential

[6] Belobaba, P.P., L.R. Weatherford. (1996). Comparing Decision Rules that Incorporate Customer Diversion in Perishable Asset Revenue Management Situations. *Decision Sciences*, Volume 27, Issue 2, pp. 343.

[7] Belobaba, P. P. (1998). Airline differential pricing for effective yield management. *The Handbook of Airline Marketing*, D. Jenkins (ed.). The Aviation Weekly Group of the McGraw-Hill Companies, New York, NY, pp. 349-361.

pricing's ability to increase revenues by offering different fare products at different prices.
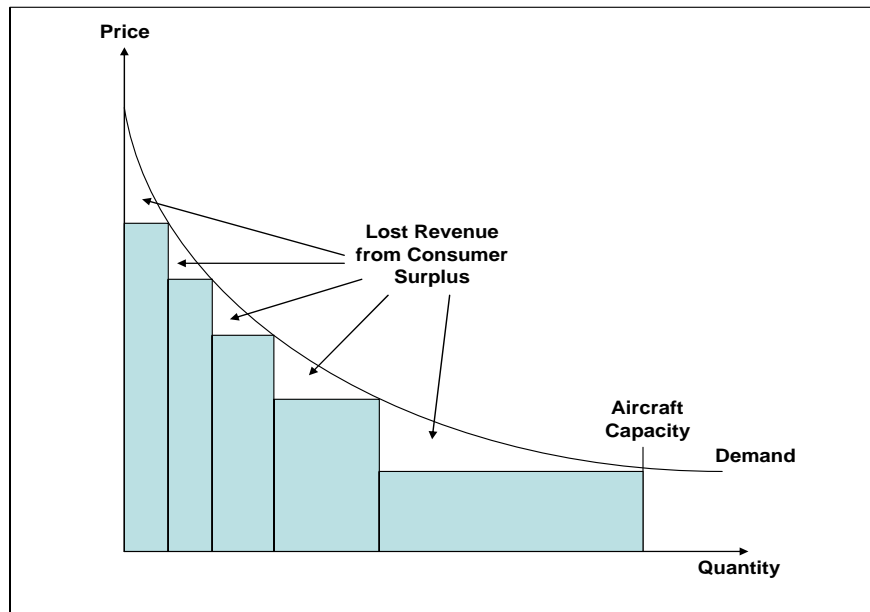
Figure 1-2: Captured Revenue Increase with Differential Pricing

The complication with differential pricing has been and continues to be the allocation of seats to the different fare classes. Since leisure travelers normally purchase tickets far in advance, airlines have to decide how many seats to allow these passengers to book and how many to withhold for later booking, higher fare business travelers. Revenue management has come a long way since 1978 in its ability to forecast demand based on multiple product offerings.

1.2 RECENT CHANGES IN THE INDUSTRY

As the airline industry moved further into the deregulated era, the practice of differential pricing became more widespread and RM systems became more adept at generating higher revenues through these practices. As airline networks grew, RM systems were allocating a common set of seats on a given flight leg between multiple fare classes across numerous itineraries. However, airlines typically operated within one fare class structure that became increasingly adept at forcing passengers to pay prices close to their WTP.

This demand segregation ability has been severely hampered by the emergence and rapid growth of LCC's. LCC's enjoy a lower cost structure than the legacy carriers, which translates into lower fares at a comparable service level and frequency. These lower fares

(and the subsequent increase in market share) have forced NLC's to match prices in markets where an LCC is present in order to maintain their market share (for more fare and seat availability matching research, refer to Lua[8]).

In addition to lower prices, LCC's have brought much more simplified (and far less differentiated) fare structures into the industry. Because of the young fleets and workforce that LCC's enjoy, they were able to offer less restricted fares in order to gain a foothold within the airline industry. The simplified fare structures themselves also helped to keep costs low as LCC's did not need expensive RM systems and departments in order to break even or even become profitable[8].

The growth of the LCC's linked with the widespread use of the Internet and the resulting transparency in ticket prices forced NLC's to match not only LCC prices, but also simplify their fare structures in these markets to compete with the less restricted fares offered by the new carriers. Delta led the NLC's with sweeping fare structure changes in 2005, with the other traditional carriers following suit soon thereafter[9]. These changes removed major restrictions, compressed the fare ratios (i.e. the ratio between the highest and lowest fare), and reduced (if not eliminated) advance purchase requirements[10].

These changes severely reduced the effectiveness of RM systems using traditional forecasting methods which operated under the assumption of demand independence. Although never entirely correct, before the emergence of LCC's there was enough differentiation in fare products that airlines had a fairly good idea which passengers were business and which were leisure. This made it easier to forecast demand and optimally allocate seats. However, the growth of simplified fare structures forced NLC's to offer a "semi-restricted" fare structure in which the lowest fares were identical except for price, while the higher fares were still differentiated according to fare restrictions.

---

[8] Lua, W.F. (2007). Matching of Lowest Fare Seat Availability in Airline Revenue Management Systems. Master's Thesis, Massachusetts Institute of Technology, Cambridge, MA.

[9] De Lollis, B. (Aug 22, 2006). Fare Changes Give Fliers Weekends Back. *USA Today*. Retrieved January 28, 2007, from the World Wide Web: http://www.usatoday.com/travel/flights/2006-08-21-saturday-stays-usat_x.htm.

[10] Delta Airlines. (Jan 5, 2005). Delta Slashes Everyday Fares up to 50 Percent as Airline Introduces SimpliFares[TM] Nationwide. Delta Airlines Press Release. Retrieved January 28, 2007, from the World Wide Web: http://news.delta.com/article9584.html.
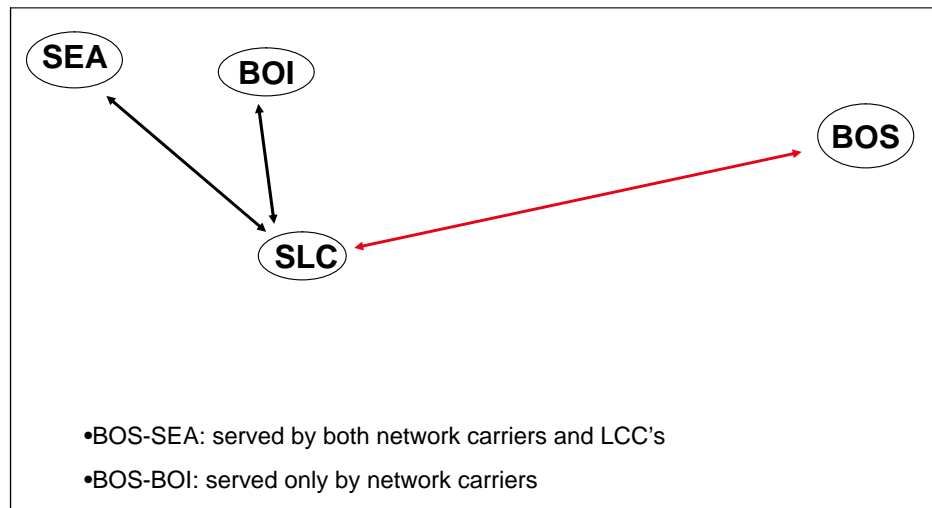
Figure 1-3: Example of Difficulties of Multiple Fare Structures for RM

This new fare structure makes the challenging problem of maximizing revenue across a network even more difficult. Figure 1-3 depicts a situation where the creation of this "semi-restricted" fare structure makes RM much more complicated for network carriers. Consider the flight leg from Boston (BOS) to Salt Lake City (SLC) on a network carrier: there will be BOS-SLC passengers on this flight, BOS-SEA passengers, Boston to Boise (BOI) passengers, as well as passengers flying from BOS and connecting in SLC to other destinations. The difficulty this new, simplified fare structure causes is twofold: first, the RM system must work within the new "semi-restricted" fare structure for the BOS-SEA passengers since an LCC is present in that market. Traditional forecasting will not be as effective in this new environment, as the demand independence assumption does not hold for the undifferentiated fares. Second, for the BOS-SLC leg, there are now two different fare structures over which the RM system must forecast and ultimately allocate seats. The BOS-SEA passengers are subject to a less restricted structure, whereas the BOS-BOI passengers book under a fully differentiated fare structure.

As this example shows, the increased complexity in the networks of legacy carriers has severely weakened the effectiveness of the traditional RM systems. Therefore, new techniques must be utilized in order to work within these multiple fare structure environments.

1.3 REVENUE MANAGEMENT DEVELOPMENTS

Given the need for better RM techniques, new methodologies for forecasting demand in the "semi-restricted" environment (Hybrid Forecasting) as well as the individual handling of multiple fare classes within the revenue management system (Fare Adjustment) will be introduced, with a more detailed examination of these techniques to follow in Chapter 3.

1.3.1 Hybrid Forecasting

In these less restricted fare structures, the RM system has great difficulty in classifying business and leisure demand as some of the classes are undifferentiated except for price, so a passenger will simply purchase the cheapest available fare. Hybrid Forecasting (HF) uses a different approach for classifying demand – "product-oriented" and "price-oriented" demand. As the names suggest, product-oriented passengers are purchasing a particular fare class because of its particular attributes (i.e. less restrictive) and not because of its price, whereas price-oriented customers are interested only in purchasing the lowest possible fare.

Since these two groups are assumed to have different booking behavior, different forecast methods are used for each group (hence the term "Hybrid Forecasting"). Belobaba and Hopperstad[11] introduced this concept so that both product- and price-oriented demand in a semi-restricted fare structure could be jointly forecasted.

The forecasting method used for the fully undifferentiated fares in these new fare structures is called Q-forecasting, also introduced by Belobaba and Hopperstad[12]. Since passengers will always buy the lowest available fare, Q-forecasting only forecasts demand at the lowest class (Q-class), then utilizes estimates of customers' WTP to forecast demand for higher fare classes (behavior known as "sell-up"). For the remaining differentiated fare classes, traditional forecasting (along with the fare class demand independence assumption) is used to forecast demand.

1.3.2 Fare Adjustment

Fare Adjustment (FA) was developed by Fiig and Isler[13] at Scandinavian Airlines (SAS) and SwissAir, respectively, to augment revenues when multiple fare structures (typically a less restricted structure and a traditionally restricted structure) exist in a network. This methodology alters the fares downward from the less restricted fare structure used in the network seat allocation optimizer, not the published fares that passengers purchase, in order to close these fare classes earlier in an attempt to force passengers to sell-up.

1.4 OBJECTIVES OF THE THESIS

Obtaining an accurate forecast for future demand is vital to the success of modern RM systems and their optimizers in particular. As the saying goes, "garbage in equals

---

[11] Belobaba, P., C. Hopperstad. (2004). Algorithms for Revenue Management in Unrestricted Fare Markets. Presented at the Meeting of the INFORMS Section on Revenue Management, Massachusetts Institute of Technology, Cambridge, MA.

[12] Belobaba and Hopperstad. (2004). "Q investigations – Algorithms for Unrestricted Fare Classes." *PODS Consortium Meeting*, Amsterdam.

[13] Fiig, T., Isler, K. (2004). "SAS O&D low cost project." *PODS Consortium Meeting*, Minneapolis, MN.

garbage out". No matter how sophisticated the RM system is, if the demand forecasts fed into the seat allocation optimizer are incorrect, the seat allocation booking limits will be erroneous as well, leading to a large deviation from the maximum revenues that could have been obtained. With multiple fare structures becoming commonplace among network carriers, the need for accurate demand forecasts is vital to maximizing revenues.

In this new environment, airlines need a way of managing each fare structure in their network separately. Fare Adjustment allows them this opportunity, as it forces the fares in the RM system from the less restricted fare structures lower, so the joint closing of fares from both fare structures can be more optimally managed.

The objective of this thesis is to evaluate the effectiveness of these two techniques in multiple fare structure environments. Hybrid Forecasting will first be evaluated by itself in these environments, and then the effect of the addition of Fare Adjustment with Hybrid Forecasting will be shown.

Previous research by Soo[14] has investigated the joint effectiveness of these two technologies, but this research was not focused on multiple fare structure environments. These techniques will be assessed in two different networks, each with 4 competing airlines: Network S1 will have a traditional fare structure and a less restricted fare structure, while Network S4 will have the same traditional fare structure but the "less-restricted" structure will be completely undifferentiated in order to fully evaluate the revenue implications of Hybrid Forecasting and Fare Adjustment. Soo's[14] analysis also concentrated on the use of arbitrary passenger sell-up rates, whereas this thesis will utilize techniques for estimating sell-up by using historical booking data.

1.5   ORGANIZATION OF THE THESIS

This thesis consists of four main parts: a literature review, description of the theory behind Hybrid Forecasting and Fare Adjustment, overview of the Passenger Origin-Destination Simulator (PODS) and the simulation environment, and the analysis from the PODS simulations.

Chapter 2 presents a summary of previous research done in Revenue Management, particularly in the areas of forecasting, sell-up estimation, specific RM techniques, and less restricted fare structures due to LCC growth.

Chapter 3 provides a detailed theoretical description of both Hybrid Forecasting and Fare Adjustment, and describes how these techniques can be beneficial in the new fare environments facing legacy carriers.

---

[14] Soo, Y.S.V. (2007). Fare Adjustment Strategies for Airline Revenue Management and Reservation Systems. Master's Thesis, Massachusetts Institute of Technology, Cambridge, MA.

The PODS simulation environment is presented in Chapter 4. Emphasis will be placed on the different network seat allocation optimizers and sell-up estimation methods used in the simulations. The two different network environments in which the simulations will be run will also be introduced.

Both Chapters 5 and 6 will contain analysis of the results from the PODS simulations. Chapter 5 will concentrate on evaluating the results of the current method of Fare Adjustment. Chapter 6 will begin with an introduction of a new Fare Adjustment methodology, and then continue with analysis from PODS using this new formulation.

The thesis concludes with a summary of experiments conducted and revenue impacts of Hybrid Forecasting and Fare Adjustment. Also presented are several opportunities for future research in this area.

CHAPTER 2
LITERATURE REVIEW

Revenue management literature dates back as far as the 1960's, with early research concentrating on the overbooking problem. However, with the reduction of regulations on the airline industry and the shrinking profit margins with which carriers have operated, much more emphasis has been placed on methods for extracting the highest possible revenues in a network, rather than maximizing the number of passengers carried. This has caused a large growth of published material from academics and practitioners alike in the past 25 years.

This chapter will consist of two sections: first, a review of revenue management techniques that have been developed for traditional fare structures. The second section will focus on the effects of Low Cost Carriers on traditional network airlines, specifically addressing the less-restricted fare structures Network Legacy Carriers are forced to offer in markets in which an LCC is present.

2.1 TRADITIONAL REVENUE MANAGEMENT MODELS

Revenue management's goal is to maximize revenues by limiting available seats to low-fare, early booking passengers in favor of high-fare, late booking passengers. Although the essential problem of calculating the correct number of seats to withhold on a given flight leg at a particular point in the booking process has been solved for certain environments, the size and complexity of current airline networks create barriers for implementation (not to mention the lack of competitive effects consideration in the algorithms). McGill and van Ryzin[4] provide a comprehensive survey of revenue management literature, while Boyd and Bilegan[15] give a more technical review of RM research. Clarke and Smith[16] review the impact of Operations Research (OR) not only on revenue management, but also on issues such as fleet assignment and infrastructure operations.

Most revenue management techniques were developed under a critical assumption that demand for a given fare class is independent of the demand for other fare classes. In the early days of revenue management, this assumption was not as detrimental to airline revenues because airlines were able to sufficiently segregate demand using fare class

---

[15] Boyd, E. A., I. C. Bilegan. (2003). "Revenue Management and E-Commerce." *Management Science,* Volume 49, Issue 10, pp.1363-1386.
[16] Clarke, M., B. Smith. (2004). Impact of Operations Research on the Evolution of the Airline Industry. *Journal of Aircraft.* Volume 41, Issue 1, pp. 62-72.

restrictions and thus target specific groups of passengers with similar sensitivities to time and price.

Barnhart et al.[17] describes the evolution of revenue management systems since the early 1980's. The earliest systems simply kept a record of bookings in the system without an effective method of using that information. These databases were upgraded so airlines could accurately track the bookings for a given flight and compare them against an expected booking curve based on previous departures of that specific flight. Currently, "third generation" RM systems contain all the previous capabilities, as well as being able to forecast demand and optimize by booking class for a future flight departure.
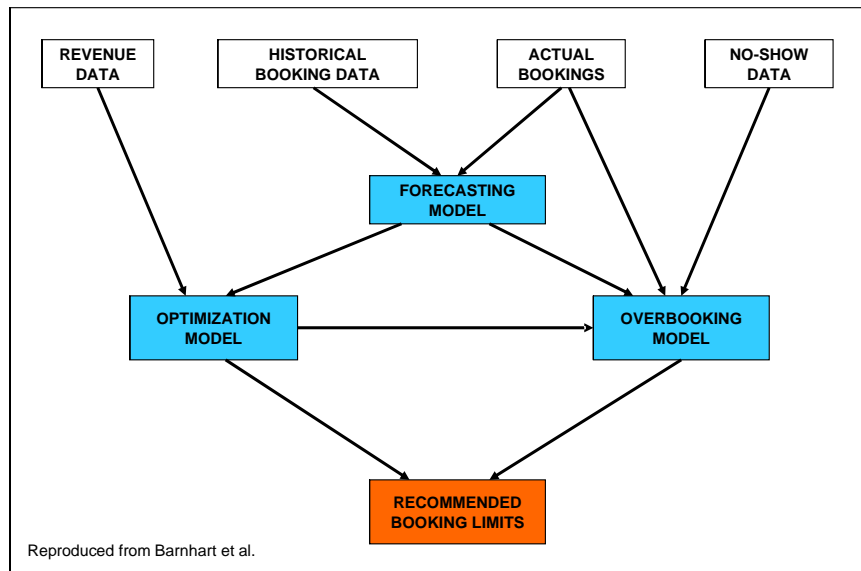
Figure 2-1: Third Generation Airline Revenue Management System

The three main components of this "third generation" RM system, as seen in Figure 2-1, are the overbooking, forecasting, and optimization models. Actual bookings for a flight leg are combined with historical booking data for the same flight leg on the same day of the week (may also be segregated by seasonality) to generate a forecast of demand for each booking class on departure day. Estimates of the revenue generated from each booking class are then combined with these demand forecasts and inputted into an optimization model which calculates the recommended booking limits for the flight. An overbooking model that uses the total demand forecasts and historical no-show rates provides the optimal overbooking level for the flight departure at the same time. Overall recommended booking limits are calculated by using the recommended booking limits from the optimization model and the optimal overbooking levels.

[17] Barnhart, C., P. P. Belobaba, A. R. Odoni. (2003). Applications of Operations Research in the Air Transport Industry. *Transportation Science*, Volume 37, Issue 4, pp. 368-391.

The following sections in this chapter will present research covering overbooking, forecasting, and seat allocation optimization.

2.1.1 Overbooking

The earliest research into the enhancement of revenue was through the use of overbooking flights. Airlines were concerned with flights departing with empty seats due to passenger no-shows and the subsequent revenue loss. Therefore, overbooking models were created in order to determine the number of bookings to accept for a flight, taking into account not only the potential for revenue loss from empty seats, but also the costs of denied boardings (passengers who have a confirmed booking on a flight but are unable to board the flight as more booked passengers arrive for the flight than the aircraft has capacity). Early deterministic overbooking models were developed by Thompson[18] and Taylor[19]. Later work by Rothstein[20,21] and Alstrup et al.[22] examined overbooking using stochastic models. In the simulation runs to be presented, no overbooking model was used.

2.1.2 Forecasting

As Figure 2-1 demonstrates, in order for a RM system to generate effective booking limits for a flight, accurate forecasts of total demand must be found by using both historical bookings on the same flight and current bookings on a flight that will depart in the future. Forecasting demand is difficult not only due to the different booking nature of passengers (i.e. low WTP customers book early, high WTP customers book late), but also because of the dynamic nature of forecasting. Forecasters must take into account how many bookings this flight has seen in the past at this time before departure and compare that with the actual bookings on the future flight at the same point in time. Large deviations from the historical booking pattern for a given flight can make it difficult on a forecaster to generate acceptable (not to mention accurate) forecasts for bookings up until flight departure.

---

[18] Thompson, H.R. (1961). Statistical Problems in Airline Reservations Control. *Operations Research Quarterly*. Volume 12, pp. 167–185.
[19] Taylor, C. J. (1962). The Determination of Passenger Booking Levels. *2nd AGIFORS Annual Symposium Proceedings*, Fregene, Italy.
[20] Rothstein, M. (1968). Stochastic Models for Airline Booking Policies. Ph.D. Thesis, Graduate School of Engineering and Science, New York University, New York, NY.
[21] Rothstein, M. (1971). An airline overbooking model. *Transportation Science*, Volume 5, pp. 180-192.
[22] Alstrup, J., S. Boaz, O.B.G. Madsen, R. Vidal, V. Victor. (1986). Booking Policy for Flights with Two Types of Passengers. European Journal of Operations Research, Volume 27, pp. 274 -288.

2.1.2.1　Pick-up Forecasting

Pick-up forecasting is a relatively intuitive way of forecasting demand at a certain point in the booking process. Instead of simply averaging the total bookings from a set of similar flights that have departed, pick-up forecasting calculates, from the historical database, the expected incremental bookings occurring in each time interval from departure. This pick-up forecast is then added to the total number of current bookings in order to forecast the total demand at the end of the time period.

There are two versions of this technique: the classical and the advanced pick-up model. The classical pick-up model only uses information from flights that have departed, whereas the advanced pick-up model developed by L'Heureux[23] also uses data from flights that have not yet departed. In this thesis, only the classical pick-up forecasting method was used. More detailed analysis of the classical pick-up forecasting model can be found in Zickus[24], Skwarek[25], Usman[26], or Gorin[27].

2.1.2.2　Other Forecasting Methods

Although the pick-up forecasting model is the only forecasting method used in this thesis, there are other RM forecasting methods that are also being used by airlines: moving average, exponential smoothing, regression, and multiplicative pick-up[28]. Zeni[29] discusses the moving average method, the multiplicative pick-up model, and exponential smoothing. Wickham[30] compared pick-up models to simple time-series and linear regression methods and found that in most cases the pick-up forecasting method led to greater revenues. Zickus[24], Skwarek[25], Usman[26], and Gorin[27] all cover different time-series models for forecasting and unconstraining historical demand.

---

[23] L'Heureux, E. (1986). A New Twist in Forecasting Short-term Passenger Pickup. *26th AGIFORS Annual Symposium Proceedings*, Bowness-on-Windemere, England, pp. 248–261.

[24] Zickus, J. S. (1998). Forecasting for Airline Network Revenue Management: Revenue and Competitive Impacts. Master's Thesis, Massachusetts Institute of Technology, Cambridge, MA.

[25] Skwarek, D. K. (1996). Competitive Impacts of Yield Management Systems Components: Forecasting and Sell-up Models. Master's Thesis, Massachusetts Institute of Technology, Cambridge, MA.

[26] Usman, A. S. (2003). Demand forecasting accuracy in airline revenue management: analysis of practical issues with forecast error reduction. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

[27] Gorin, T. O. (2000). Airline revenue management: sell-up and forecasting algorithms. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

[28] Weatherford, L. (1999). Forecast Aggregation and Disaggregation. *IATA Revenue Management Conference Proceedings*.

[29] Zeni, R. H. (2001). Improved Forecast Accuracy in Revenue Management by Unconstraining Demand Estimates from Censored Data. Ph.D. Thesis. Rutgers, the State University of New Jersey, Newark, NJ.

[30] Wickham, R. R. (1995). Evaluation of Forecasting Techniques for Short-Term Demand of Air Transportation. Master's Thesis, Massachusetts Institute of Technology, Cambridge, MA.

2.1.3 Seat Allocation Optimizers

The ultimate goal of RM, deciding how many seats to allocate to different fare classes in order to maximize revenue, is calculated by the seat allocation optimizer using the demand forecasts and revenue data. There are two different categories of algorithms used to set the booking limits: leg base control and Origin-Destination (OD) control.

2.1.3.1   Leg-Based Control

After receiving demand forecasts for individual fare classes, leg-based optimizers determine the availability of fare classes on each leg. One of the seminal works of RM, written by Littlewood[31], solved the fare class mix allocation problem for two classes by introducing an idea known as "nesting". Instead of allocating a certain number of seats to each fare class independently (and allowing a lower fare class to be available when a higher fare class is not), nesting protects higher fare classes from lower fare classes by limiting the number of seats sold in lower fare classes according to the forecasted demand and expected seat revenue for each fare class.

Belobaba[32], in his Ph.D. thesis, extended the nested seat allocation problem to multiple fare classes using the Expected Marginal Seat Revenue (EMSR) heuristic. EMSR is simply the average fare of the seat being considered multiplied by the probability that demand will materialize for that marginal seat. This seat should only be held for a particular fare class when its EMSR is greater than the average fare of the next lower class. Therefore, the booking limit for a fare class is found when the EMSR for that fare class equals the average fare of the next lower class. Belobaba[33] further updated this methodology in order to make it more robust, and it has now become known as the EMSRb method. More in-depth information on the EMSRb algorithm can be found in Williamson[34] and Lee[35].

Curry[36], Wollmer[37], and Brumelle and McGill[38] have formulated optimal nested seat allocation algorithms for multiple fare classes. However, the EMSRb method continues

[31] Littlewood, K. (1972). Forecasting and Control of Passenger Bookings. *12thAGIFORS Annual Symposium Proceedings*, Nathanya, Israel, pp. 95–117.

[32] Belobaba, P. P. (1987). Air Travel Demand and Airline Seat Inventory Management. Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA.

[33] Belobaba, P. P. (1992). The Revenue Enhancement Potential of Airline Revenue Management Systems. *ASTAIR Proc. Adv. Software Tech. Air Transport,* London, U.K.

[34] Williamson, E. L. (1992). Airline Network Seat Inventory Control: Methodologies and Revenue Impacts. Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA.

[35] Lee, A. Y. (1998). Investigation of Competitive Impacts of Origin-Destination Control using PODS. Master's Thesis, Massachusetts Institute of Technology, Cambridge, MA.

[36] Curry, R. E. (1990). Optimal Airline Seat Allocation with Fare Classes Nested by Origin and Destinations. *Transportation Science*. Volume 24, Issue 3, pp. 193–204.

[37] Wollmer, R. D. (1992). An Airline Seat Management Model for a Single Leg Route when Lower Fare Classes Book First. *Operations Research.* Volume 40, Issue 1, pp. 26–37.

to be much more widely used by airlines throughout the world because of its short computation time and its results that continue to be shown close to optimal.

2.1.3.2   Origin-Destination Control

Although leg-based RM systems increase revenues, there is still a disconnect between the way these systems forecast and optimize, by leg, and the way passengers book air travel, by path (which may consist of more than one leg). Leg-based systems are able to maximize yield but not maximize revenue in the network. In order for a passenger to book on a connecting itinerary, the same fare class must be available on all legs of the itinerary. Therefore, high demand legs which shut down lower fare classes may cause connecting passengers with higher total revenue contributions to be rejected in favor of local passengers.

In an effort to better align RM methods to the booking process as well as maximize total network revenues, Origin-Destination revenue management systems were developed. These systems allocate seat inventory based on paths, not merely on individual legs, making them valuable to airlines that operate vast hub-and-spoke networks.

Smith and Penn[39] at American Airlines developed one of the first methods for OD control called "virtual buckets". Instead of fare classes, an Origin-Destination Fare (OD Fare) is placed into a virtual bucket (which is internal to the airline's RM system) that allows both local and connecting itineraries to be compared. Booking limits are then set on these virtual buckets rather than on actual fare classes. However, this is a "greedy" approach that favors connecting passengers even in situations where two local passengers would contribute more revenue than a connecting passenger.

In an effort to better account for the bias toward connecting passengers, Displacement Adjusted Virtual Nesting[39] (DAVN) was created. This method corrects the value of the OD Fare for the displacement costs of taking a seat from a local passenger. A deterministic linear program (LP) is run to find leg shadow prices, which are estimates of the leg displacement costs[40]. DAVN then controls inventories based on the Network Revenue Value or Pseudo Fare, which is the original OD Fare minus the leg displacement costs over a particular path. Williamson[41], Vinod[42], and Wei[43] offer more information on both OD control and specifically the DAVN approach.

---

[38] Brumelle, S. L. and McGill, J. I. (1988). Airline Seat Allocation with Multiple Nested Fare Classes. Paper presented at the Fall ORSA/TIMS Conference, Denver, CO. Also presented at the University of British Columbia, 1987.

[39] Smith, B. C., C. W., Penn. (1988). Analysis of Alternative Origin-Destination Control Strategies, *AGIFORS Symposium Proceedings*, 28, 123-144.

[40] Belobaba, P. (2008). "Overview of PODS Simulation Models and RM Systems." *PODS Consortium Meeting*, Los Angeles, CA.

[41] Williamson, E. L. (1988). Comparison of optimization techniques for origin-destination seat inventory control. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

Another approach to OD control that is much simpler to implement is "Bid-Price Control". Developed by Simpson[44] and Williamson[45], and also discussed by Smith and Penn39, this method only makes the airline store a bid-price value (approximated displacement cost) for each leg in the network. An OD Fare is then compared to the sum of the bid prices on the legs that the path crosses. If the OD Fare is greater than this itinerary bid-price, the booking is accepted, otherwise it is rejected.

A number of algorithms for calculating leg bid prices have been developed, including Network Bid-Price (NetBP), Heuristic Bid-Price (HBP) by Belobaba[46], and Probabilistic Bid-Price (ProBP) by Bratu[47].

2.2   RISE OF THE LOW COST CARRIER

Although Southwest began operations in 1971[48], the deregulation of the US airline industry destroyed many of the barriers that had previously discouraged new airlines, especially with a LCC business model, from entering into the lucrative US domestic market. Deregulation in other world markets has also caused an influx of LCC's, some of whom have become quite large and profitable airlines. Tretheway[49] writes in his 2004 article that in the US, Canada, and European Union, an LCC carrier has the highest market capitalization of any airline (Southwest, WestJet, and Ryanair, respectively).

LCC's have been able to gain market share quickly through their low cost structures. These airlines are relatively new, which means operating newer aircraft and younger workforces. The labor cost advantage has been one of the main reasons LCC's have been able to offer low fares for extended periods of time (however, NLC's have done a very good job of lower labor costs in recent years, due in no small part to bankruptcy negotiations).

---

[42] Vinod, B. (1995). Origin and Destination Yield Management. *The Handbook of Airline Economics*, D. Jenkins (ed.). The Aviation Weekly Group of the McGraw-Hill Companies, New York, NY, 459-468.

[43] Wei, Y. J. (1997). Airline O-D Control using Network Displacement Concepts. Master's Thesis, Massachusetts Institute of Technology, Cambridge, MA.

[44] Simpson, R. W. (1989). Using Network Flow Techniques to Find Shadow Prices for Market and Seat Inventory Control, Memorandum M89-1, MIT Flight Transportation Laboratory, Massachusetts Institute of Technology, Cambridge, MA.

[45] Williamson, E. L. (1992). Airline Network Seat Inventory Control: Methodologies and Revenue Impacts. Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA.

[46] Belobaba, P. P. (1998). The evolution of airline yield management: fare class to origin-destination seat inventory control. *The Handbook of Airline Marketing,* D. Jenkins (ed.). The Aviation Weekly Group of the McGraw-Hill Companies, New York, NY, pp. 285-302.

[47] Bratu, S. J-C. (1998). Network value concept in airline revenue management. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

[48] Gittell, J. H. (2003). *The Southwest Airlines way: using the power of relationships to achieve high performance.* McGraw-Hill, New York, NY.

[49] Tretheway, M. W. (2004). Distortions of airline revenues: why the network airline business model is broken. *Journal of Air Transport Management*. Volume 10, Issue 1, pp. 3-14.

LCC's have typically focused on the price-oriented passenger, reducing the need for high service amenities such as first-class seats and business lounges in airports[50]. NLC's have realized the advantages of limiting the number of customer types they attempt to serve in order to better provide services to all passengers. Although some legacy carriers have undoubtedly turned their focus toward the price-oriented passenger, Elliott[51] shows anecdotal evidence that United Airlines has instead decided to focus on the product-oriented customer.

Another problem NLC's are encountering is the growing price sensitivity of product-oriented passengers[52,53]. Some companies are no longer willing to pay the last minute, full-fare ticket on NLC's anymore. Not only are substitutes such as video teleconferencing gaining ground, but LCC's have noticed this change in passenger behavior and have acted upon it. Southwest and JetBlue have recently began offering amenities and fare products that are intended to cater to the more price-sensitive business traveler[54].

Gorin's[55] thesis focuses on the impact of the entrance of LCC's into the US airline industry, and both Weber and Thiel[56] and Dunleavy and Westerman[57] provide good comparisons between the business models of LCC's and NLC's.

2.2.1 Less-Restricted Fare Structures

One of the main characteristics of LCC's is their simpler, less-restricted fare structures. In contrast to the legacy carriers who offer many different fare products with many different restrictions and price levels, LCC's tend to only have a few fare products which are priced low and have fewer restrictions. This fare structure does not allow sophisticated RM systems to allocate seats based on segregated demand, but LCC's use this to their advantage, as these simplified structures do not need large RM and OR

---

[50] Shumsky, R. (2006). The Southwest effect, airline alliances and revenue management. *Journal of Revenue and Pricing Management*, Volume 5, Issue 1, pp. 83-89.

[51] Elliott, C. (Jan 17, 2006). When Fliers Benefit From Airline Bankruptcy. *New York Times*. Retrieved February 19, 2008, from the World Wide Web: http://www.nytimes.com/2006/01/17/business/17soff.html.

[52] Bender, A. R., F. J. Stephenson. (1998). Contemporary Issues Affecting the Demand for Business Air Travel in the United States, *Journal of Air Transport Management*, Volume 4, Issue 2, pp. 99-109.

[53] Cary, D. (2004). Future of Revenue Management: A view from the inside. *Journal of Revenue and Pricing Management,* Volume 3, Issue 2, pp. 200-203.

[54] McCartney, S. (Feb 19, 2008). Unusual Route: Discount Airlines Woo Business Set. *Wall Street Journal*. Retrieved February 19, 2008, from the World Wide Web: http://online.wsj.com/article/SB120338017328875635.html?mod=dist_smartbrief.

[55] Gorin, T. O. (2004). Assessing low-fare entry in airline markets: impacts of revenue management and network flows. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA.

[56] Weber, K., Thiel, R. (2004). "Optimisation issues in low cost revenue management." *AGIFORS Reservations & Revenue Management Study Group Meeting*, Auckland, New Zealand.

[57] Dunleavy, H., Westermann, D. (2005). Future of airline revenue management. *Journal of Revenue and Pricing Management*. Volume 3, Issue 4, pp. 380-282.

departments and allow the airline to keep costs low. According to Tretheway[49], "the introduction of low one way fares…has undermined the price discrimination ability of the FSNC's [full service network carrier], and is the most important pricing development in the industry in the past 25 years."

With these new, low-priced fare products being offered, legacy carriers were forced to match the LCC's on markets in which they were directly competing with an LCC or else risk losing a significant portion of market share. However, the NLC's revenue management systems were still operating under the assumptions of segmented and independent demand. When the legacy carriers switched to these less-restricted fare structures, they lost the ability to effectively segment and forecast demand.

2.2.1.1 Spiral Down Effect

With these segmentation restrictions removed, NLC's were no longer able to force product-oriented passengers to buy in higher fare classes. This allowed product-oriented passengers to "buy-down" – purchasing a ticket in a fare class lower than what they would have previously purchased because they can now pay less for the same product. This diversion of high-fare demand to lower fare classes can lead to a cycle known as "spiral down". Figure 2-2 illustrates this revenue-lowering cycle.
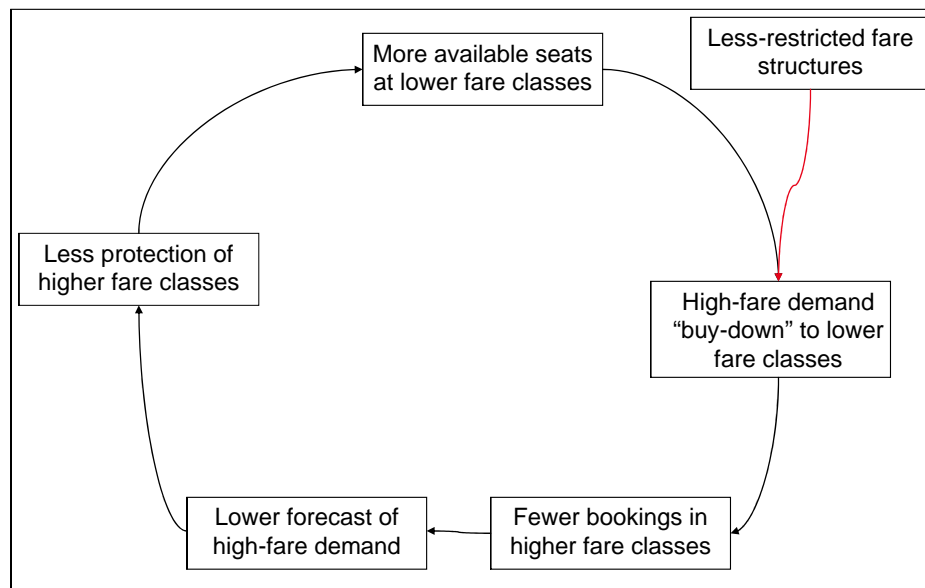
Figure 2-2: Spiral Down Effect[8]

The less-restricted fare structures reduce the ability to segment demand, and therefore high-fare demand buys-down into the lower fare classes. The historical booking database receives fewer bookings in the higher fare classes which causes the revenue management system to forecast a lower amount of high-fare demand. This forecast is then fed into the seat allocation optimizer, which proceeds to protect fewer seats for higher fare classes since the demand forecast is lower. In turn, this makes more seats available for lower

booking classes, which encourages even more high-class demand buy-down. Obviously, this cycle will repeat with more and more seats being opened for lower fare classes and revenues being constantly diluted.

A mathematical model to describe the spiral down effect is developed by Cooper et al.[58], while the effects of this cycle have been analyzed both by practitioners at United Airlines[59] and by Cusano[60] with PODS.

2.2.2 Fare and Seat-Availability Matching

Although many NLC's have invested millions of dollars into sophisticated RM systems, many times they make decisions based on the competition. Although the RM system may have a particular fare class opened or closed, an analyst may manually override the system if a competitor's actions jeopardize the carrier's market share. Nason[61] points out that the pricing transparency brought on by the online travel agencies has had a large effect on consumer behavior and will have a large effect on the future of revenue management.

Both Nason[61] and Cary[53] argue that in the future, revenue gains may be possible by incorporating real-time competitor information into an airline's revenue management system. One airline, bmi, has actually began incorporating this technology into its RM system by scanning other airline fares, comparing the fares with bmi's by flight/date, and integrating this information with bmi's RM system[62]. However, this is done at an "*ad hoc* individual level", and the information is not currently integrated into the seat allocation optimizer for bmi, but rather used by an analyst to make more well-informed decisions.

Although fare matching has allowed NLC's to compete based on price, they are still limited by their cost structures[63]. An NLC may not be able to operate for a prolonged period while matching an LCC if its cost structure is too high. However, as noted above, the costs for NLC's and LCC's have been converging in recent years as LCC work forces and equipment age and NLC's have gained concessions at the bargaining table.

---

[58] Cooper, W. L., T. Homem-de-Mello, A. J. Kleywegt. (2004). Models of the Spiral-down Effect in Revenue Management. Working Paper, Department of Mechanical Engineering, University of Minnesota, Minneapolis, MN.

[59] Ozdaryal, K., Saranathan, B. (2004). "Revenue management in broken fare fence environment." *AGIFORS Reservations & Revenue Management Study Group Meeting*, Auckland, New Zealand.

[60] Cusano, A. J. (2003). Airline revenue management under alternative fare structures. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

[61] Nason, S.D. (2006). Forecasting the future of airline revenue management. *Journal of Revenue and Pricing Management*, Volume 4, Issue 1, pp. 64-66.

[62] Donnelly, S., James, A., Binnion, C. (2004). bmi's response to the changing European airline marketplace. *Journal of Revenue and Pricing Management*. Volume 3, Issue 1, pp. 10-17.

[63] Forsyth, P. (2003). Low-cost Carriers in Australia: Experiences and Impacts. *Journal of Air Transport Management*, Volume 9, Issue 5, pp. 277-284.

In Lua's[8] thesis, he investigates the impacts of an airline matching the fare seat availability of another airline in a competitive environment using PODS. He concluded that in most simulations, the airline that instituted fare seat availability matching lost revenue while the airline that was matched improved its revenues.

There has been a large amount of research in the past 25 years on all aspects of revenue management. However, the growth of the LCC's has invalidated some assumptions and has rendered some advances moot. The next chapter will describe in detail two different methods for maximizing revenue in less differentiated multiple fare structure environments.

2.3   CHAPTER SUMMARY

This chapter began with a review of the literature on the three traditional revenue management models, with the discussion focused largely on forecasting and seat allocation optimization. Next, the literature on the emergence of LCC's was explored. The impact of LCC's on legacy carriers has been immense, as the removal of fare class restrictions and subsequent demand segmentation degeneration has led to lower revenues for the NLC's. The traditional revenue management systems employed by the NLC's are ill-suited for use on less restricted fare structures, as section 2.2.1.1 explained. Finally, a short review of fare and seat-availability matching was conducted as the Internet has had a large effect on the actions of not only consumers but also RM analysts as more information is readily available.

CHAPTER 3
DEVELOPMENT OF NEW RM METHODS FOR SIMPLIFIED FARE STRUCTURES

The introduction of less-restricted fare structures by LCC's has had a large impact on the revenue maximizing ability of NLC's. In order to retain market share, NLC's have been forced to match the fare structures of LCC's in markets where an LCC is present. This presents a challenge for legacy carriers as their traditional revenue management techniques will result in spiral down in these new fare structures. Additionally, a legacy carrier uses not one but several different fare structures across its network. This introduces added complexity to the RM problem as a single flight leg can carry passengers from both traditional and less restricted fare structures. Hybrid Forecasting (HF) and Fare Adjustment (FA) are two techniques used to help the NLC's recapture some of the lost revenue due to the simplification of fare structures and removal of segmentation restrictions.

3.1   HYBRID FORECASTING

In order to match LCC's in a given market, legacy carriers have had to create semi-restricted fare structures. These structures are characterized by lower fare classes that have the same set of restrictions and are distinguished only by price, and higher fare classes which still segregate demand effectively through different combinations of restrictions. HF utilizes a new classification for passengers in the forecasting process: yieldable (product-oriented) and priceable (price-oriented)[64].

There will be two subsequent sections: first, a new forecasting technique solely for price-oriented passengers will be covered, and second, the incorporation of this new forecasting method with traditional pick-up forecasting for product-oriented passengers will be discussed, resulting in Hybrid Forecasting.

3.1.1 Q-Forecasting

Since the traditional assumption of demand independence is no longer valid for undifferentiated fares, a new forecasting approach was developed by Belobaba and Hopperstad[65] called Q-forecasting. This method assumes fully undifferentiated fares such that passengers buy the lowest available fare.

---

[64] Boyd, E. A., Kallesen, R. (2004). The science of revenue management when passengers purchase the lowest available fare. *Journal of Revenue and Pricing Management*. Volume 3, Issue 2, pp. 171-177.
[65] Belobaba, P., C. Hopperstad. (2004). Algorithms for revenue management in unrestricted fare markets. Presented at the Meeting of the INFORMS Section on Revenue Management, Massachusetts Institute of Technology, Cambridge, MA.

Q-forecasting's goal is to force sell-up by estimating a passenger's WTP and closing down lower fare classes so that the passenger must purchase a fare closer to their WTP. Instead of forecasting demand for each class separately (as is done with independent fare classes), Q-forecasting only forecasts demand at the lowest fare class (denoted class Q) and then uses estimates of the customers' WTP to force sell-up by closing down lower fare classes. An overview of the Q-booking process is presented in Figure 3-1.
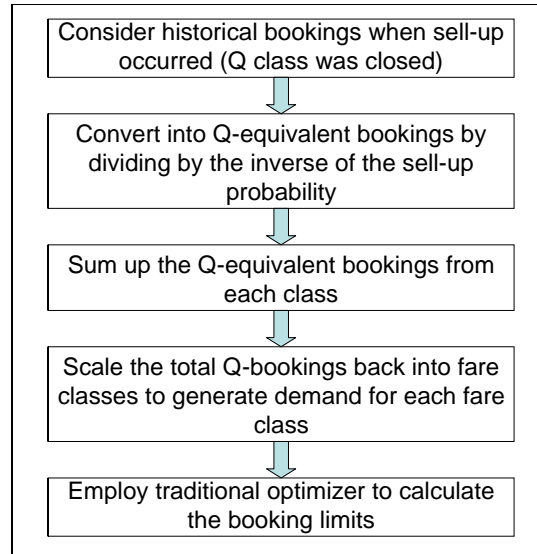
Figure 3-1: Basic Overview of Q-forecasting[66]

First, historical bookings made when sell-up occurred (Q class was closed) are converted into Q-equivalent bookings by scaling them by the inverse of the sell-up probability:

$$\text{Total Q - equivalent bookings} = \sum_{f} \frac{\text{Historical bookings in class } f}{\text{Probability of sell-up from } Q \text{ to } f}$$

An example of the calculation of Q-equivalent bookings is shown in Table 3-1.

Class f	Historical bookings in f	Probability of sell-up from Q to f	Q-equivalent bookings for class f
Y	5	15%	5/.15 = 33
B	15	40%	15/.40 = 38
M	25	80%	25/.80 = 31
Q	45	100%	45

Total Q-equivalent bookings = 147

Table 3-1: Example Calculation of Q-equivalent Bookings

[66] Cléaz-Savoyen, R. L. (2005). Airline Revenue Management Methods for Less Restricted Fare Structures. Master's Thesis, Massachusetts Institute of Technology, Cambridge, MA.

Using the sell-up probability from Q to some higher class to scale the bookings is, in effect, calculating the number of bookings that would have been seen had class Q been open.

Once the Q-equivalent bookings have been calculated for each class, they are all summed together. Finally, this total Q-booking number is then repartitioned back into individual fare classes by multiplying the total Q-bookings by the probability that a passenger will sell-up to class f but not $f - 1$:

$$\text{Forecasted demand for class } f = \text{Total Q - equivalent bookings} \cdot [(\text{Probability of sell - up from}$$
$$Q \text{ to } f) - (\text{Probability of sell - up from } Q \text{ to } f - 1)]$$

Using the values from Table 3-1 (total Q-equivalent bookings = 147), Table 3-2 shows the repartitioning of demand back to each individual fare class.

Class f	Probability of sell-up from Q to f	Repartitioned demand for class f
Y	15%	147 * (.15-.00) = 22
B	40%	147 * (.40 - .15) = 38
M	80%	147 * (.80 - .40) = 59
Q	100%	147 * (1.00 - .80) = 29

Table 3-2: Example Calculation of Repartitioning Demand to each Fare Class

Cléaz-Savoyen[66] showed in his thesis that the use of Q-forecasting on unrestricted fare structures is effective in reducing the revenue loss due to the reduction of fare class restrictions.

Although Q-forecasting overcomes the independence assumption, it requires willingness-to-pay inputs (sell-up). This has become a significant stumbling block as effective methods for estimating sell-up have not performed as well as arbitrarily inputted sell-up values in the Passenger Origin-Destination Simulator. In the next chapter, two sell-up estimation methods will be presented that will be used in the simulations.

3.1.2 Combining Forecasting Techniques

Q-forecasting deals only with the undifferentiated fares in the semi-restricted fare structure. It is not valid to use Q-forecasting for the entire semi-restricted fare structure as bookings will be made in a higher fare class when it is not the lowest open class (which would violate the assumption of Q-forecasting that all classes are undifferentiated). However, using traditional pick-up forecasting for the entire fare structure is also not correct with the lower fare classes violating the demand independence assumption.

Initially suggested by Boyd and Kallesen[64], Belobaba and Hopperstad[65] developed Hybrid Forecasting in order to forecast both price- and product-oriented demand simultaneously for a semi-restricted fare structure. The lower fare classes will have demand forecasts created solely by price-oriented passengers, while higher fare classes will have their demand forecasts created by a combination of price- and product-oriented passengers. Reyes finds in his thesis that the use of hybrid forecasting on semi-restricted fare structures gives approximately 3% revenue gains over using standard pick-up forecasting[67].

3.2   FARE ADJUSTMENT

Fare Adjustment (FA) was developed by Fiig and Isler[68] as a way to de-couple fare structures that could exist on a single leg. FA allows the airline to independently control the more-restricted fare structure and the less-restricted fare structure on the leg. Section 1-2 describes a situation where FA could be used to manage the fare structures on the BOS-SLC leg. The following sections will explain the need for FA and demonstrate the method for calculating the adjusted fare.

3.2.1 Overview of the Need for Fare Adjustment

Fare Adjustment was created for use with DAVN, where a pseudo fare (OD Fare – displacement costs, section 2.1.3.2) is mapped to a virtual class within the RM system, and the booking limits are set on these virtual classes. It is possible for a legacy airline to have an OD fare class from a more-restricted fare structure and a separate OD fare class from a less-restricted fare structure to be mapped to the same virtual bucket. Demand for the more-restricted fare structure will be independent as it is effectively segregated, whereas demand for the less-restricted structures will be dependent and revenues are maximized by estimating sell-up for both the forecaster and the optimizer. Therefore, it may be optimal to close down a virtual class for one fare structure, but not optimal for the other fare structure, and vice versa. An example of this conflict is shown in Figure 3-2. In this example, if virtual class 4 is shut down, both of the pseudo fares mapped to that bucket would be shut down, which may not be optimal overall for the network.

---

[67] Reyes, M. H. (2006). Hybrid Forecasting for Airline Revenue Management in Semi-Restricted Fare Structures. Master's Thesis, Massachusetts Institute of Technology, Cambridge, MA.
[68] Isler and Fiig, (2004). SAS O&D Low Cost Project, PODS presentation, Minneapolis-St Paul.
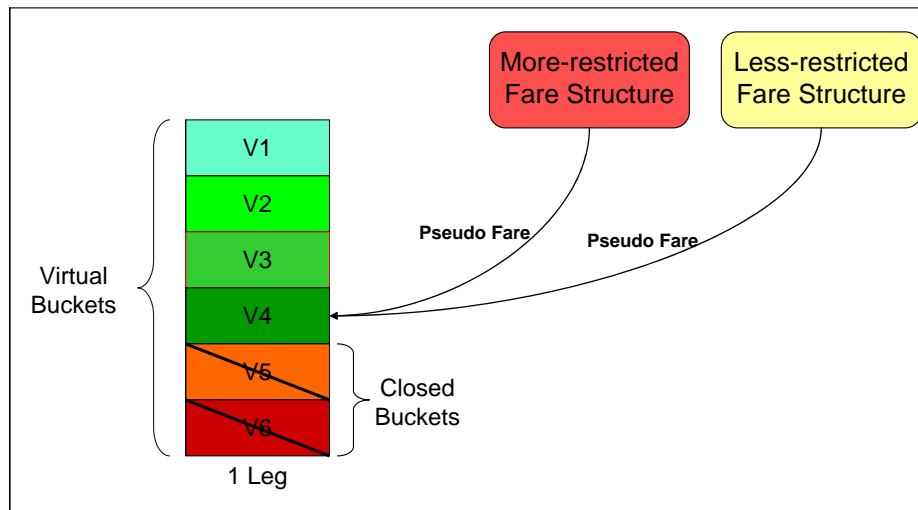
Figure 3-2: Example of Multiple Fare Structure Conflict on a Leg[66]

It is beneficial to quickly review DAVN in order to more fully understand the implementation of Fare Adjustment. DAVN is an Origin-Destination optimizer because it uses path/class forecasting instead of simply leg-based forecasting. Origin-Destination Fares (OD Fares) are inputted into a network LP which calculates the displacement cost for a given leg. This displacement cost is a penalty applied to connecting itineraries for the possible displacement of a local passenger. Although initially working with OD Fares, the actual value that is bucketed and optimized is the pseudo fare. For a given itinerary, the pseudo fare for a given leg on that itinerary is the OD Fare minus the displacement cost of all other legs on that itinerary. Therefore, the pseudo fare for a local passenger is simply the passenger's OD Fare, while a connecting passenger's pseudo fare is the difference between the actual fare on the leg the passenger is booking and the displacement costs of the other legs in the itinerary. Although DAVN uses path/class demand forecasts, all the fare classes sharing a particular leg are bucketed in terms of pseudo fares and optimized, meaning the seat allocation is calculated on a leg level.

The Fare Adjustment method, instead of feeding the network LP with OD Fares, gives the LP the Marginal Revenue (MR, calculation shown later in this chapter), or adjusted fare, of an itinerary. This means the OD Fare has another value, the Price-Elasticity Cost (PE Cost), subtracted from it to account for the risk of buy-down, therefore "Marginal Revenue = OD Fare – PE Cost". The LP then calculates the displacement cost of each leg based on the inputted Marginal Revenue values for each OD fare class on the network. Whereas the value bucketed without FA was the pseudo fare (OD Fare – Displacement costs), the bucketed value with FA is "OD Fare – Displacement costs – PE Cost". The inclusion of MR into the DAVN process is shown in Figure 3-3.
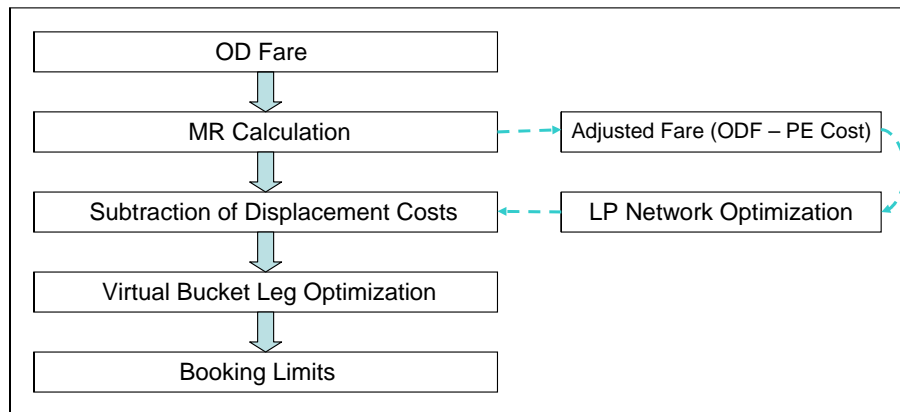
Figure 3-3: Addition of Marginal Revenue into the DAVN Process[66]

The subtraction of this PE Cost causes the undifferentiated fares to be mapped to a lower bucket and allows the two different fare structures to be managed more separately, as shown in Figure 3-4. Note that differentiated fare classes are not affected by the PE Cost as there is no risk of buy-down in effectively segmented demand.
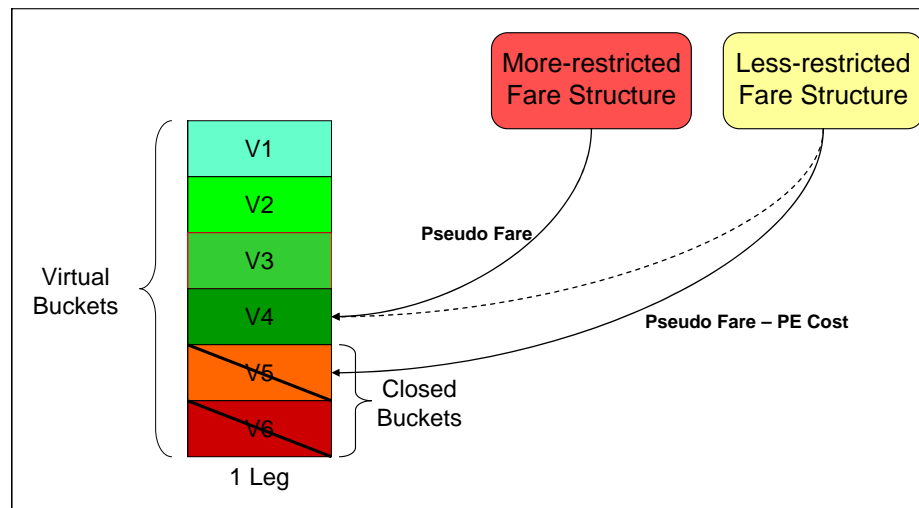
Figure 3-4: Example of Decoupling Multiple Fare Structures[66]

The following sections will show how to calculate the adjusted fare through the Marginal Revenue Transformation[69].

---

[69] Fiig, T., K. Isler, C. Hopperstad, P. Belobaba, (2008). Optimization in Airline Fare Structures with Restricted and Unrestricted Products. Draft Paper. PODS Consortium.

3.2.2 Introduction of the Optimization Problem

In order to develop the Fare Adjustment theory, a simple, general problem will be introduced. This problem will be the basis for an example that will demonstrate the validity of the Marginal Revenue Transformation. This general problem is a single flight leg with a set capacity (CAP). The time dimension will not be considered as this will be a static optimization. Additionally, the problem will be made more simplistic as demand is assumed to be deterministic.

The airline can have any number of fare classes C_i with $i = 1..n$ with 1 being the highest class. The fares associated with each class are in decreasing order with fare class 1 having the highest fare. The objective of the airline is to choose a strategy S_i (which set of classes to open) that maximizes revenue. For the general problem and the following example, all fare classes are nested, such that strategy S_4 has fare classes 1-4 open.

For the general problem, demand for an individual fare class is dependent on the availability of the other fare classes, or put another way, on the optimization strategy chosen. Therefore, demand for a given fare class is denoted $d_i(S_k)$. Using the demand categorization method suggested by Boyd and Kallesen, price-oriented demand for a fare class is given as d_i^{price} and product-oriented demand as d_i^{prod}. Finally, total demand for a chosen strategy is calculated as the sum of the demands for the open fare classes and denoted Q_i, while total revenue is simply the demand for a fare class multiplied by its fare, summed over all the available fare classes and denoted TR_i.

Each of the notations described is given in Table 3-3, which shows the demand by class, total demand and total revenue for each optimization strategy.

Set	C_1	C_2	C_3	...	C_n	Dem.	TR
S_1	$d_1(S_1)$					Q_1	TR_1
S_2	$d_1(S_2)$	$d_2(S_2)$				Q_2	TR_2
S_3	$d_1(S_3)$	$d_2(S_3)$	$d_3(S_3)$			Q_3	TR_3
...						...	...
S_n	$d_1(S_n)$	$d_2(S_n)$	$d_3(S_n)$	...	$d_n(S_n)$	Q_n	TR_n

Table 3-3: Notation for Demands and Total Revenue by Strategy[69]

For the remaining sections in this chapter, the following example will be used to illustrate the usefulness of the Marginal Revenue Transformation methodology. The aircraft capacity is CAP=100. The flight leg under consideration will have a fare structure that consists of four fare classes with the corresponding prices and demands listed in Table 3-4.

Fare Class	Demand (d_i)	Fare Price (f_i)
1	20	1000
2	18	650
3	35	400
4	58	275

Table 3-4: Demands and Fare Prices for Four-Class Example

3.2.2.1   Fully Restricted Fare Structure

As previously discussed, a fully restricted fare structure effectively segments demand by fare class. Due to these restrictions, passengers book in a particular fare class regardless of the availability of other fare classes. In this structure, the demand for any fare class is no longer $d_i(S_k)$ but rather d_i, as demand does not depend on the strategy chosen. Since there is no sell-up or buy-down in a fully restricted structure, there is no demand classified as price-oriented, so demand for a fare class is simply d_i^{prod}. Total demand Q_i is the sum of the individual fare class demands from 1 to i, and total revenue is calculated $TR_i = \sum_{j=1,...,i} d_j f_j$.

It is evident that the optimal seat allocation policy for a fully restricted fare structure is to make seats available from the highest fare class downward until either demand is satisfied or capacity is reached. Table 3-5 shows that the first three fare classes receive a number of seats equal to demand, while fare class 4 can only be allocated 27 seats as capacity is reached. This is the optimal allocation for this deterministic example, giving a total revenue of $53,125.

	Fare	Demand Booked	Revenue
1	1000	20 20	20,000
2	650	18 18	11,700
3	400	35 35	14,000
4	275	58 27	7,425

Total Revenue = $53,125

Table 3-5: Bookings with Demands and Corresponding Revenues for Fully Restricted Fare Structure

40

3.2.2.2   Fully Unrestricted Fare Structure

With the absence of any restrictions or advance purchase requirements, a fully unrestricted fare structure has no demand segmentation ability, and the only difference between fare classes is the price. As opposed to the previous section where bookings were made in multiple fare classes, all passengers who are willing to pay more than the price of the lowest offered fare will simply buy-down and book in the lowest available fare class.

With this fare structure, demand for a fare class is affected by the availability of other fare classes as the demand independence assumption is no longer valid. Whereas demand for a fare class in the previous section was d_i, a fully unrestricted structure makes the optimization strategy a factor, so fare class demand is given as $d_i(S_k)$. However, with no restrictions, 100% buy-down into the lowest available class will occur. Therefore, all fare classes that are not the lowest available fare class will have zero demand. Since no passengers will book based on the attributes of a fare class, all of the demand for the lowest available class is price-oriented and denoted d_i^{price}.

In order to optimize a fully unrestricted fare structure, the lowest available class to be left open must be determined in order to maximize revenue. Because of the 100% buy-down, demand for a class C_i, given it is the lowest available class, is the sum of the product demands for C_i and all higher classes. For this example, $d_2^{price} = d_1 + d_2 = 38$.

An equivalent way of determining the price-oriented demand for a fare class is using the sell-up rate $psup_i = Q_i / Q_n$, the ratio of the total demand when class i is the lowest open to the total demand of leaving all classes open. A more intuitive explanation of sell-up is the probability a passenger, given he is willing to book in class n, will book in class i if it is the lowest available class. Therefore, the demand for the lowest available class can also be calculated $d_i(S_i) = Q_n psup_i$.

As Table 3-6 illustrates, the optimal policy for this example is to close class 4 and have class 3 be the lowest available class. Notice that the plane is not filled to capacity (73 passengers), but in implementing the strategy S_4 to fill every seat, total revenue would have been lost as passengers would have booked at $275 instead of $400. There is also a large difference in the total revenues of the optimal strategies for the two different fare structures as the loss of demand segregation leads to large revenue losses.

	Fare	Demand Booked	Revenue
1	1000	20 0	20,000
2	650	38 0	24,700
3	400	73 73	29,200
4	275	131 0	27,500

<div style="text-align:right; color:red">Total Revenue = $29,200</div>

Table 3-6: Bookings with Demands and Corresponding Revenues for Fully Unrestricted Fare Structure

3.2.3 Marginal Revenue Transformation

In section 3.2.1, it was shown that Fare Adjustment changes the value inputted to the LP, and that instead of inputting the OD Fare, a value called the PE Cost is subtracted from the OD Fare to account for the possibility of buy-down and the LP calculates the displacement cost of each leg using these pseudo fares. The altered input value to the LP is also called the Marginal Revenue for an itinerary. If the variable costs of producing an extra airline seat are ignored, the Marginal Revenue methodology says to "allocate seats according to their marginal revenue until the marginal revenue of the last seat is negative or capacity is reached"[69]. The following sections will introduce the Marginal Revenue Transformation and prove its validity through application to the preceding examples.

This methodology is very helpful for airline RM systems, as it allows a general problem with dependent demand to be essentially transformed into independent demand. The next sections will show that if we know the optimal strategy for a problem with dependent demand and perform the Marginal Revenue Transformation on the fare classes, then treat the demand for these new classes as independent, the same optimal strategy will again be found. Then, since the transformed fare classes are treated as independent, traditional leg and OD optimizers that operate under the demand independence assumption can be used.

The Marginal Revenue Transformation takes the original fare classes and maps them into primed booking classes $C_k^{'}$ with corresponding primed demands and fares as shown. Since the fare classes are assumed to be nested, the primed demand for a fare class k is the difference in total demands when k is the lowest available class and when $k-1$ is the lowest available. The primed fare is the marginal revenue for that fare class. It calculates the change in total revenue given the change in total demand if fare class k is opened.

MARGINAL REVENUE TRANSFORMATION[69]:

$$d_k^{'} = Q_k - Q_{k-1}, k = 1,...,n .$$

$$f_k^{'} = \left(TR_k - TR_{k-1}\right) / \left(Q_k - Q_{k-1}\right), k = 1,...,n$$

The marginal revenue value for a fare class is used to determine the optimal strategy by determining which fare classes should be open or closed. Since the methodology states that seats should only be allocated to fare classes with positive marginal revenues, negative marginal revenue fare classes will be closed, and the optimal strategy will be obtained.

For the Marginal Revenue Transformation methodology to be valid, two things must be shown: first, that the primed fare classes are indeed independent, and second, that treating these primed fare classes as independent yields the same demand and revenue as the dependent structure for any given strategy. Fiig et al. provide analysis where both arguments are proven true, therefore the Marginal Revenue Transformation can be applied to the preceding examples to show that the optimal strategy is identical using either methodology.

3.2.4 Application of Marginal Revenue Transformation

In the following subsections, the original single leg, deterministic demand optimization problem that was solved earlier will be re-solved using the Marginal Revenue Transformation. This will show for a very simplified example that the optimal strategy is the same when using the transformation to create independent fare classes. The notation used will be identical to that introduced in the preceding sections.

3.2.4.1 Fully Restricted Fare Structure

The Marginal Revenue Transformation can be applied to any general fare structure in order to create primed fare classes that are independent, essentially creating a fully restricted fare structure. Obviously, a fare structure that is already fully restricted does not need to be transformed, or in the same way, applying the Marginal Revenue Transformation to a fully restricted fare structure will generate the exact same fare structure.

The marginal revenue value that is fed into the LP is calculated "OD Fare – PE Cost", and the PE Cost is only applied when there is a risk of buy-down. Put another way, the PE Cost is only applicable to fare classes with price-oriented demand. For a fully restricted fare structure, there is no risk of buy-down and consequently no PE Cost for any fare class. Therefore, the full OD Fare value is fed into the LP, just as the Marginal Revenue Transformation shows that the adjusted fare equals the OD Fare.

3.2.4.2   Fully Unrestricted Fare Structure

Since all demand in a fully unrestricted fare structure is classified price-oriented, each fare class (except for the highest fare class because it will be open until capacity is reached) has an associated PE Cost and subsequently a marginal revenue value that is lower than the current fare. The mapping of the unrestricted fare classes to independent primed fare classes is done as outlined in section 3.2.3 (general solutions in Fiig et al.).

C_i	f_i	Q_i	TR_i	$\Delta TR_i = TR_i - TR_{i-1}$	$d_i^{'} = \Delta Q = Q_i - Q_{i-1}$	$f_i^{'} = \Delta TR / \Delta Q$
1	1000	20	20,000	$(20{,}000 - 0) =$ 20,000	$(20 - 0) =$ 20	1,000
2	650	38	24,700	$(24{,}700 - 20{,}000) =$ 4,700	$(38 - 20) =$ 18	261.11
3	400	73	29,200	$(29{,}200 - 24{,}700) =$ 4,500	$(73 - 38) =$ 35	128.57
4	275	131	27,500	$(27{,}500 - 29{,}200) =$ -1,700	$(131 - 73) = 58$	-29.31

Table 3-7: Marginal Revenue Transformation on Fully Unrestricted Structure

Table 3-7 shows the Marginal Revenue Transformation on the example used in this chapter. The first three fare classes have positive marginal revenues and would be available for bookings. The negative marginal revenue for the lowest fare class indicates that if this class were made available, total revenue would decrease from the preceding strategy of fare class 3 being the lowest available class. Clearly, fare class 4 would not have any seats allocated to it and would be closed. This is the identical strategy that was calculated in section 3.2.2.2.

The table shows that the fare for the highest fare class remains unchanged, while the adjusted fare for fare class 2 is 261.11, a substantial reduction from the actual fare of $650. Fare class 4's marginal revenue is negative, which means that if this class were opened, the revenue lost from passengers who would have booked in class 3 booking in class 4 would have been greater than the revenue generated from the increased number of bookings. This is more easily explained by examining fare class 2.

When fare class 2 is opened, revenue is lost due to the buy-down of passengers from class 1 to 2. These 20 passengers, instead of paying $1000, are paying $650, a loss in revenue of 20(1000-650) = $7000. In contrast to the fully restricted fare structure when the incremental revenue of opening class 2 is 18*650 = $11,700 because no buy-down occurs, the incremental revenue for class 2 in the fully unrestricted fare structure is reduced by this buy-down revenue loss, making the marginal revenue
18*650-$7000=$4700. Consequently, the marginal revenue for fare class 2 is 4700/18 = 261.11.
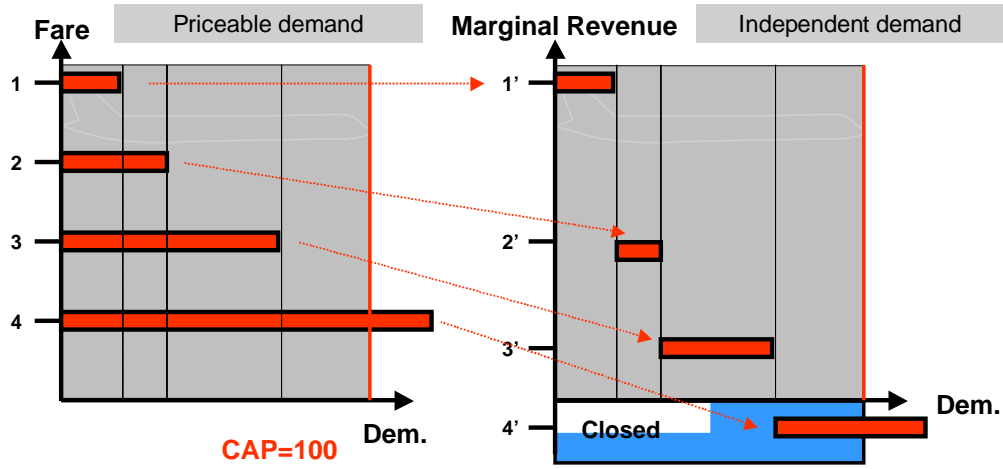
Figure 3-5: Mapping of Unrestricted Fare Classes to Independent (Primed) Fare Classes[69]

Figure 3-5 shows the transformation of fully unrestricted fare classes into independent, primed fare classes. Table 3-7 shows the marginal revenue values for the primed classes on the right. Even without the actual marginal revenue figures, the optimal strategy could be found from a graph such as Figure 3-5, as any fare class mapped to a primed fare class with negative marginal revenue should be closed.

3.2.4.3   Semi-Restricted (Hybrid) Fare Structure

Finally, the Marginal Revenue Transformation can be used on a semi-restricted fare structure, which is much more common in the airline industry, particularly for legacy airlines.   In this structure, some classes will be fully unrestricted with the only differentiator being price, while other classes will be fully restricted with independent demand.

When a hybrid fare structure is applied to the current example, demand for a fare class cannot be wholly classified either price- or product-oriented, but instead is a combination of both categories of demand. It is evident, then, that the total demand for a fare class is the product-oriented demand along with the demand that will sell-up to the given class but not to the next higher class, $d_k^{'} = d_k^{prod} + Q_n\left(psup_k - psup_{k-1}\right)$[69].

Fiig et al. introduce a new variable x_k defined as the ratio of product-oriented demand to the total demand for a fare class. Solving for the marginal revenue in terms of x_k gives $f_k^{'} = x f_k + (1-x) f_k^{'}$ (refer to Fiig et al. for derivation). It can be seen that the marginal revenue for a fare class in a hybrid fare structure is the unadjusted fare weighted by the product-oriented demand and the adjusted fare weighted by the price-oriented demand.

Obviously, classes with a larger proportion of price-oriented demand will have lower adjusted fares and therefore be more likely to be closed.

The hybrid fare structure marginal revenue equation is important because it provides a link between restrictions in a fare structure and the aggressiveness of the optimization strategy. The less the demand segmentation ability, the higher the risk of buy-down, and therefore the marginal revenue of the fare classes will be lowered further to mitigate that risk.

To illustrate Fare Adjustment for a semi-restricted fare structure, the following four-class flight leg example in Table 3-8 will be used (note the fares and demands are different than the previous example):

C_i	f_i	d_i^{price}	d_i^{prod}	$psup_i$
1	1000	5	15	10%
2	750	12	11	35%
3	450	32	15	70%
4	175	77	0	100%

Table 3-8: Fares, Demands and Sell-Up Rates for Semi-Restricted Fare Structure Example

The price-oriented demand for each fare class is the number of price-oriented bookings that occur in that fare class if it is the lowest available class. By definition, the lowest possible fare class cannot have any product-oriented demand.

As seen earlier, a fully restricted fare structure does not need to adjust fares as there is no risk of buy-down. Since there is only buy-down risk for fully unrestricted fare structures, the final adjusted fare for a fare class in a hybrid fare structure will be lower than the actual fare only if that fare class has some price-oriented demand associated with it. Therefore, the adjusted fare for the price-oriented demand is calculated as before, then it is combined with the ratio of product-oriented demand and unadjusted fare to find the final adjusted fare for a fare class. Table 3-9 illustrates the Marginal Revenue Transformation performed on the price-oriented demand, and Table 3-10 shows the final adjusted fares.

C_i	f_i	d_i^{price}	TR_i^{price}	$\Delta TR_i^{price} = TR_i^{price} - TR_{i-1}^{price}$	$d_i^{'price} = \Delta d_i^{price} = d_i^{'price} - d_{i-1}^{'price}$	$f_i^{'price} = \dfrac{\Delta TR_i^{price}}{\Delta d_i^{price}}$
1	1000	5	5,000	(5,000 – 0) = 5,000	(5 – 0) = 5	1,000
2	750	12	9,000	(9,000 – 5,000) = 4,000	(12 – 5) = 7	571.43
3	450	32	14,400	(14,400 – 9,000) = 5,400	(32 – 12) = 20	270.00
4	175	77	13,475	(13,475 – 14,400) = -925	(77 – 32) = 45	-20.56

Table 3-9: Marginal Revenue Transformation of Price-Oriented Demand in a Hybrid Fare Structure

C_i	$d_i^{'} = d_i^{prod} + Q_n(psup_i - psup_{i-1})$	$f_i^{'} = x f_i + (1-x) f'_i$
1	$15 + 77 \cdot (0.10 - 0.0) = 23$	$(\dfrac{15}{20}) \cdot 1000 + (1 - \dfrac{15}{20}) \cdot 1000 = 1000$
2	$11 + 77 \cdot (0.35 - 0.10) = 30$	$(\dfrac{11}{23}) \cdot 750 + (1 - \dfrac{11}{23}) \cdot 571.43 = 656.45$
3	$15 + 77 \cdot (0.70 - 0.35) = 42$	$(\dfrac{15}{47}) \cdot 450 + (1 - \dfrac{15}{47}) \cdot 270.00 = 327.47$
4	$0 + 77 \cdot (1.0 - 0.70) = 23$	$(\dfrac{0}{77}) \cdot 175 + (1 - \dfrac{0}{77}) \cdot -20.56 = -20.56$

Table 3-10: Total Marginal Revenue Transformation of a Hybrid Fare Structure

This example shows the connection between restrictions and aggressiveness of fare adjustment. With more restrictions and a greater ability to segment demand, most of the demand is product-oriented, thereby weighting the unadjusted fare more heavily and reducing the risk that a particular fare class is shut down. Conversely, in fare structures that can only weakly segment demand, the adjusted fare from the price-oriented demand is more heavily weighted, driving the final adjusted fare lower and increasing the likelihood a fare class is closed in order to account for the buy-down possibility.

This chapter has reviewed two new RM methods designed to reduce the revenue loss of simplified and mixed fare structures: Hybrid Forecasting and Fare Adjustment. The next chapter will cover the Passenger Origin-Destination Simulator (PODS) which will be used for the evaluation of the performance of these two new RM methods. The implementation of HF and FA in PODS will also be shown, as well as the mechanism in PODS that links the aggressiveness of the Fare Adjustment to the corresponding fare structure.

3.3 CHAPTER SUMMARY

In this chapter, two new techniques for reducing the revenue loss from less restricted fare structures were introduced: Hybrid Forecasting and Fare Adjustment. The two forecasting techniques for the two different types of demand were presented, as Q-forecasting is used for price-oriented demand, while standard pick-up forecasting continues to be used for product-oriented demand. These two demand forecasts are then combined in HF to create a single forecast value for each fare class.

Next, Fare Adjustment was discussed in detail. First, the need for FA was developed using the virtual buckets of DAVN as an example of how two different fare structures could become coupled and lead to sub-optimal decisions. Fiig et al.'s Marginal Revenue Transformation was then presented through a common example applied to different fare structures. In a semi-restricted fare structure, the adjusted fare is a combination of the full, unadjusted fare and the fully adjusted fare, weighted by the product and price demand forecasted for the fare class. Finally, a new example was given for a semi-restricted fare structure to illustrate FA's ability to close classes when there is a risk of buy-down.

CHAPTER 4
PODS SIMULATION ENVIRONMENT

Simulation is often better suited to model the competitive airline environment as static analytic revenue management models must assume away certain parts of the problem in the name of simplification. Although this allows for rigorous treatment of different revenue management techniques, competitive behavior between airlines and the passenger choice dynamic are usually excluded from the analysis[70]. Simulation-based analysis allows for revenue management practices to be modeled in a dynamic environment complete with the interactions between passenger choice and the revenue management system.

This chapter first presents an overview of the Passenger Origin-Destination Simulator (PODS) which is used to test both Hybrid Forecasting and Fare Adjustment. This overview will detail the components of the passenger choice model, the revenue management system, and the interaction between these two elements. The main part of the chapter will be devoted to the implementation of different seat-allocation optimizers, sell-up estimators, and the two new RM methods in PODS. Finally, the simulated four-airline environments used for analysis will be introduced.

4.1 PODS ARCHITECTURE

PODS, which was originally developed by C. Hopperstad, M. Berge, and S. Filipowki at the Boeing Company, was developed from Boeing's previous Decision Window Model (DWM)[71] which determined a passenger's choice based on multiple variables, such as schedules and airline characteristics. However, the DWM left out the fares and restrictions associated with these fares, factors that would also influence passenger choice. Although the bulk of the passenger choice model in PODS replicates the DWM, a passenger's choice set now includes multiple fare products. The incorporation of fare products into the passenger choice model allow the effects of a new revenue management technique to be tested in a competitive environment.

A PODS simulation is referred to as a "run". A "run" consists of both "trials" and "samples": a run of our four-airline environment will consist of two trials, and a single trial consists of 600 samples. A sample is a single departure day (i.e. a Friday), and this day is replicated 600 times in order to ensure statistically significant results. The overall operating statistics for a "run" are the average of the two "trials".

---

[70] Gorin, T., P. Belobaba. (2004). Revenue management performance in a low-fare airline environment: insights from the Passenger Origin-Destination Simulator. *Journal of Revenue and Pricing Management.* Volume 3, Issue 3, pp. 215-236.
[71] Boeing Airplane Company. (1997). Decision Window Path Preference Methodology Description. Seattle, WA.

In order to start a PODS run, user-defined inputs must be fed into the system. As the simulation progresses, these values are eventually replaced by calculated values. In order to analyze results generated by the simulation, the first 200 samples of every trial are discarded. Therefore, the results of a 600-sample trial are only based on the last 400 samples, and a full PODS simulation run is the averaged result of 800 samples.

For each sample, the booking process in PODS begins 63 days before departure and is broken into 16 succeeding time frames that end on the departure day. At the beginning of the booking process the time frames last a week, but as the departure day nears the length of the time frames decrease in anticipation of increased booking activity as shown in Table 4-1. The revenue management systems of the airlines update the path/class availabilities at the start of each time frame, while passenger events such as bookings and cancellations occur randomly within each time frame.

Time Frame	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
Days to Departure	63	56	49	42	35	31	28	24	21	17	14	10	7	5	3	1	0

Table 4-1: Booking Process Time Frames

PODS is a simulator that links together passenger choice with airline RM systems in order to analyze the effectiveness of different RM techniques. The third generation RM system (see section 2.1) calculates the air travel options to offer to prospective passengers. This information is then passed to passengers searching for air travel in a particular OD market who then choose a particular airline, path, and fare class based on the options available from the RM system and the passenger's characteristics. The PODS structure is shown in Figure 4-1.
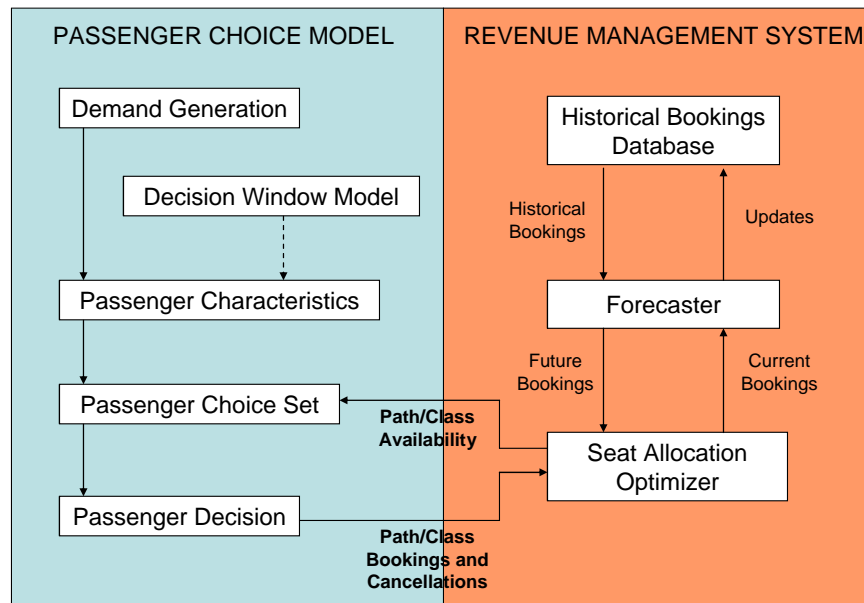
Figure 4-1: PODS Structure

4.1.1 Passenger Choice Model

The effectiveness of a RM technique and its ability to maximize revenues is dependent on the response (or lack thereof) from passengers. The Passenger Choice Model in PODS manages the behavior of passengers and their choices relative to path/class availability through four succeeding steps: Demand Generation, allocation of Passenger Characteristics, identification of a Passenger Choice Set, and ultimately the Passenger Decision. This section will provide an overview of the Passenger Choice Model, but a more in-depth discussion with model validation can be found in Carrier[72].

4.1.1.1   Demand Generation

In this step, average daily air travel demand based upon data from the PODS Consortium airline members is generated for every OD market in the network. This total generated demand is then partitioned between leisure and business passengers. Variability is then generated randomly around this average daily demand, creating the daily demand curve for each type of passenger. However, seasonality and day-of-week variability is not included in the demand generation process. Lastly, the arrival pattern through the booking process for each group of passengers is modeled according to user-defined booking curves: the booking curves used in this thesis for both business and leisure passengers are shown in Figure 4-2. These curves are also based on data from the PODS Consortium's airline members that show leisure travelers tend to arrive earlier in the booking process than business travelers.

---

[72] Carrier, E. (2003). Modeling Airline Passenger Choice: Passenger Preference for Schedule in the Passenger Origin-Destination Simulator (PODS). Master's Thesis, Massachusetts Institute of Technology, Cambridge, MA.
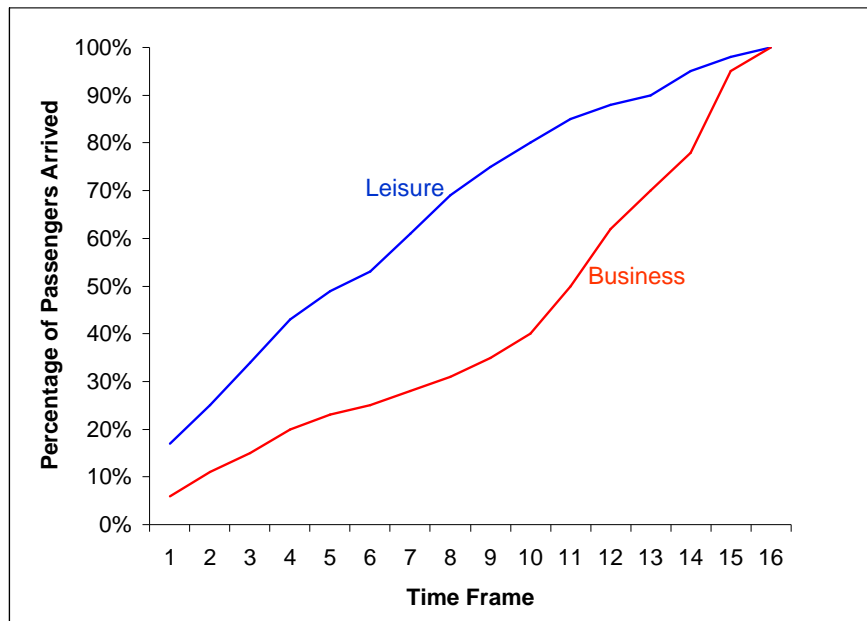
Figure 4-2: PODS Booking Curves by Passenger Type

The demand generation process in PODS is impacted by two inputs: the base fares and the number of passengers willing to travel at those base fares. For a given base fare, there is an associated number of passengers of each passenger type willing to travel at that fare. In PODS, the user can generate higher or lower levels of demand by manipulating a setting called the Demand Multiplier (DM). The average demand resulting from the base fare and passenger numbers is given a DM of 1.00. Different multiples of the baseline demand can be obtained by scaling this number. In this thesis, Fare Adjustment will be tested in the average demand environment of 1.00 and lower demand environments using Demand Multipliers of 0.90 and 0.80.

4.1.1.2   Passenger Characteristics

In the second step of the Passenger Choice Model, three sets of characteristics are assigned to the passengers generated by the preceding step: a decision window, a maximum willingness-to-pay (WTP), and a set of disutility costs indicating the passenger's sensitivity to different aspects of a possible booking.

Each passenger's decision window is defined by the earliest acceptable departure time and the latest acceptable arrival time. Business travelers are given smaller decision windows than leisure travelers demonstrating their time sensitivity. At this point, without factoring in WTP and disutility costs, all path and fare class combinations which are fully contained within a passenger's decision window are equally appealing, while those not fully within the decision window are all equally unappealing as they require re-planning of the decision window.

Next, the passenger is given a maximum WTP, the maximum amount a passenger is willing to pay for air travel. The WTP is calculated from price-demand curves for both business and leisure travelers, which are calculated with user inputs from the formula:

$$P(\text{pay at least } f) = \min[1, e^{\frac{\ln(0.5) \cdot (f - basefare)}{(emult - 1) \cdot basefare}}]$$

where:
f is the fare;
basefare is an input in PODS, at which an input specified number of passengers are willing to pay to travel;
emult is the elasticity multiplier such that 50% of passengers are willing to pay emult*basefare to travel;

Leisure travelers are given a smaller *emult*, which results in a quick drop off in their price-demand curve to reflect their high price sensitivity. Any fare that exceeds a passenger's maximum WTP will be removed from that passenger's choice set.

Finally, a passenger is assigned disutility costs that are randomly generated from a probability distribution based on his passenger type. These disutility costs represent the passenger's sensitivity to schedule preference (re-planning for a path outside the decision window), path quality (non-stop versus connecting itineraries), and restrictions associated with a fare product (Saturday night stay, non-refundability, change fee). A more comprehensive discussion of disutility costs in PODS can be found in Lee[73].

4.1.1.3   Passenger Choice Set

After the passenger has been assigned a full set of characteristics, he is presented with a set of paths and fare classes from which to choose (the passenger always has the "do-nothing" option as well). These travel options are provided from the Seat Allocation Optimizer in the Revenue Management System (see Figure 4-1). Some of the options will immediately be removed from the passenger's choice set, as either the RM system of one or more airlines has closed down a path/fare class in the desired OD market, advance purchase requirements cannot be met, or the fare is higher than the passenger's maximum WTP.

4.1.1.4   Passenger Decision

In order to make a decision, the passenger will sum up the fare and relevant disutility costs of each available option (called the total generalized cost) and choose the option with the lowest generalized cost (do-nothing or "no-go" alternative included). This booking is then fed back into the airline's RM system as the available seat inventory is decreased and the historical bookings database is increased.

---

[73] Lee, S. (2000). Modeling Passenger Disutilities in Airline Revenue Management Simulation. Master's Thesis, Massachusetts of Technology, Cambridge, MA.

4.1.2 Revenue Management System

The airline side of PODS consists of a third-generation RM system that is made up of three components: a Historical Bookings Database, a Forecaster, and a Seat Allocation Optimizer. The link between the passenger side and the airline side in PODS exists between the Seat Allocation Optimizer of the RM system and the Passenger Choice Set and Passenger Decision of the Passenger Choice Model as seen in Figure 4-1.

4.1.2.1   Historical Bookings Database

In PODS, the Historical Bookings Database records the fare class and path of every booking on a particular airline. Initially it is filled with default bookings provided by the user, but as the simulation progresses those are eventually substituted with actual bookings from the simulation (hence the first 200 samples are burned in a PODS trial). The user defines the number of observations of a given flight held in the database. For the simulations in this thesis, 26 samples (previous departures) of each flight are held in the Historical Bookings Database to be used by the forecaster.

4.1.2.2   Forecaster

The forecaster in PODS takes booking data directly from the Historical Bookings Database in order to provide a forecast of future demand by fare class and path. As described in section 3.1, pick-up forecasting for product-oriented demand and Q-forecasting for price-oriented demand is combined to provide a path/class forecast through Hybrid Forecasting.

The data taken directly from the Historical Bookings Database is inherently biased because it only reports demand that actually booked an itinerary. There may have been prospective passengers who wished to book travel but were unable due to the seat availability of a fare class. Therefore, the demand data given to the forecaster must be "unconstrained", which means estimating the number of bookings that would have occurred had the particular fare class been open indefinitely. In this thesis, the Booking Curve detruncation method, a percentage-based multiplier that extrapolates demand for closed classes from trend data from open classes, is used. More detailed analysis of detruncation techniques can be found in Usman[26] and Gorin[27].

4.1.2.3   Seat Allocation Optimizer

With the capacities of flight legs in the upcoming PODS runs ranging from 50 to 150, the Seat Allocation Optimizer must allocate these perishable seats in such a way as to try and maximize revenue. Each airline in a simulation run can be given a different Seat Allocation Optimizer, and PODS allows the user to choose from a number of different optimizers with a range of sophistication levels. The three main optimizers used in thesis will be Adaptive Threshold 90 (AT90), EMSRb, and DAVN. The results of two other optimizers will be shown for comparison purposes in the next chapter, Heuristic Bid-

Price (HBP) and Probabilistic Bid-Price (PBP). Descriptions of these two optimizers can be found in Cléaz-Savoyen[66].

4.1.2.3.a   First Come First Served (FCFS)

This inventory control mechanism actually performs no optimization at all, instead allowing all passengers to book in any fare class that has not been closed due to advance purchase requirements until capacity is reached. This is normally used as a way of measuring the performance of RM systems.

4.1.2.3.b   Threshold Algorithms

The first step to improving upon the FCFS method and actually managing booking limits is by using a Threshold Algorithm. A load factor threshold, between 0% and 100%, is set for each class and when the threshold is met, the class is closed down. There are two ways of implementing a Threshold Algorithm: Fixed Threshold and Adaptive Threshold. The Fixed Threshold, as it indicates, uses load factor thresholds inputted at the beginning of the simulation and stays constant all the way through the booking process. An Adaptive Threshold Algorithm has an overall target load factor – in this thesis 90% - set at the onset of the simulation, and the individual fare class thresholds are computed at each time frame by the bookings received up to that point in the booking process in order to achieve the target (90%) load factor. The user can also control the fluctuations of the fare class thresholds by assigning minimum and maximum bounds on the individual load factors (see Lua[8]). In this thesis, this Seat Allocation Optimizer is used to imitate a simple revenue management system employed by a LCC.

4.1.2.3.c   Fare Class Yield Management (FCYM)

This leg-based method sets booking limits on nested fare classes based on the Expected Marginal Seat Revenue (EMSRb[33], see section 2.1.3.1), the expected amount of revenue to be obtained by making the next seat available for a given class. Therefore, EMSR = OD Fare * probability of selling the seat. Under the assumption of demand independence, an incremental seat is held for a given fare class as long its EMSR is greater than the EMSR of the class beneath it. An example of booking limits set by using the FCYM method is shown in Table 4-2.

Class	Fare	Avg demand	Stdev demand	Booking Limit
1	450	18	6	*100*
2	325	21	7	*86*
3	200	28	9	*62*
4	125	35	11	*30*

Table 4-2: Example of EMSRb with Capacity = 100

The booking limit for each class is the maximum number of seats that the airline should sell in that particular fare class. For example, if the booking limit for fare class 3 is reached, any subsequent request for a seat will only be presented with fare class 2 or 1 as available booking options. The nested booking limits for this example are shown in Figure 4-3.
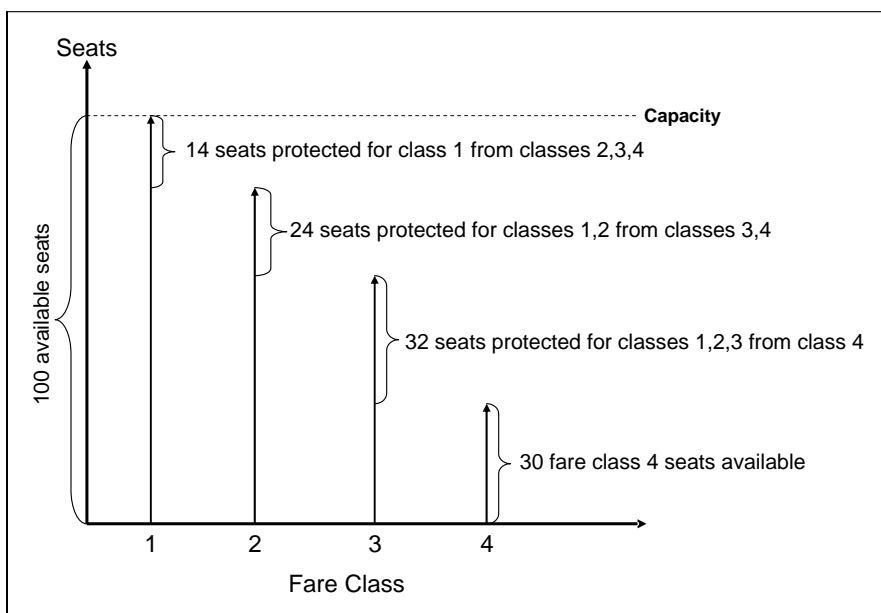
Figure 4-3: Nested Booking Limits with Capacity = 100

4.1.2.3.d   Displacement Adjusted Virtual Nesting (DAVN)

DAVN, as described in both sections 2.1.3.2 and 3.2.1, utilizes path/class forecasting in a network LP to calculate the displacement costs of a connecting passenger for each leg. Williamson[34] provides more detail on the network LP problem. The displacement costs are then subtracted from the OD Fare for a particular itinerary and mapped into a virtual bucket, which is then managed at a leg level.

Since Fare Adjustment was designed for use with this method, simulation runs with DAVN will be more heavily analyzed when determining the effectiveness of Fare Adjustment.

4.2 SELL-UP ESTIMATION

The Revenue Management system uses a passenger's maximum willingness-to-pay (WTP) in order to calculate the probability the passenger will sell-up to a higher fare class than he otherwise would have booked. This sell-up probability is used in both Q-forecasting (section 3.1.1) and the EMSRb algorithm (discussed in more detail in Cléaz-Savoyen[66]) to force passengers to pay closer to their max WTP's. This section will

discuss the sell-up implementation in PODS as well as two methods for estimating sell-up from historical bookings.

4.2.1 Sell-Up Implementation in PODS

The probability of sell-up in PODS is based on a negative exponential distribution. Therefore, the probability a random passenger would book in class f given that he would have booked in the lowest class (denoted Q) and no other lower class is available is calculated:

$$psup_{Q \to f} = e^{-(\frac{fare_f}{fare_Q}-1) \cdot econ_{tf}}$$

where: $fare_f$ is the price of the higher fare class;

       $fare_Q$ is the price of the lowest fare class;

       $econ_{tf}$ is the sell-up constant for time frame tf ;

In order to calculate the sell-up probability to a fare class f, the airline must input a value for $econ_{tf}$.

$$econ_{tf} = -\frac{\ln(\frac{1}{2})}{Frat5_{tf} - 1}$$

Frat5 is the fare ratio from the lowest fare class at which 50% of the passengers will sell-up to a higher fare class. Frat5 is used by airlines in PODS to determine passengers' willingness-to-pay, so a higher frat5 means that the fare ratio where 50% of the passengers will sell-up to is higher, corresponding to lower price sensitivity, higher WTP's and higher probabilities of sell-up.

Since it is assumed that business travelers (who also have higher WTP's) tend to arrive later in the booking process than leisure passengers, the frat5 curve in PODS generally increases following an "S-curve" toward the departure date.
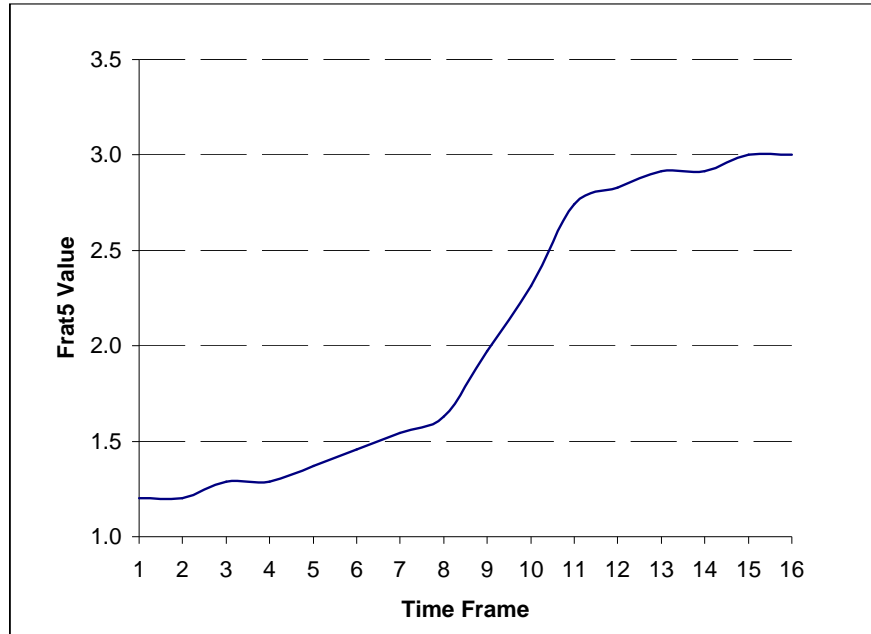
Figure 4-4: Example of Frat5 S-curve (Frat5C)

As Figure 4-4 shows, the S-shape of the Frat5 curve captures the increase in WTP throughout the booking process. At 63 days from departure (time frame 1), 50% of the passengers would only be willing to sell-up to a fare of $120 (assuming the base fare is $100). However, with only one day before departure (time frame 16), 50% of the passengers would be willing to sell-up to a fare of $300. This particular frat5 curve is denoted Frat5C as created by Cléaz-Savoyen[66], who created five arbitrary frat5 curves of varying aggressiveness for sensitivity analysis.

When an airline in PODS uses a higher frat5 curve to calculate the sell-up probabilities, it is assuming its passengers have a high WTP and thus will protect more seats for higher fare classes in order to force passengers to book in those fare classes. Figure 4-5 illustrates this higher probability of sell-up with three frat5 curves with increasing aggressiveness.
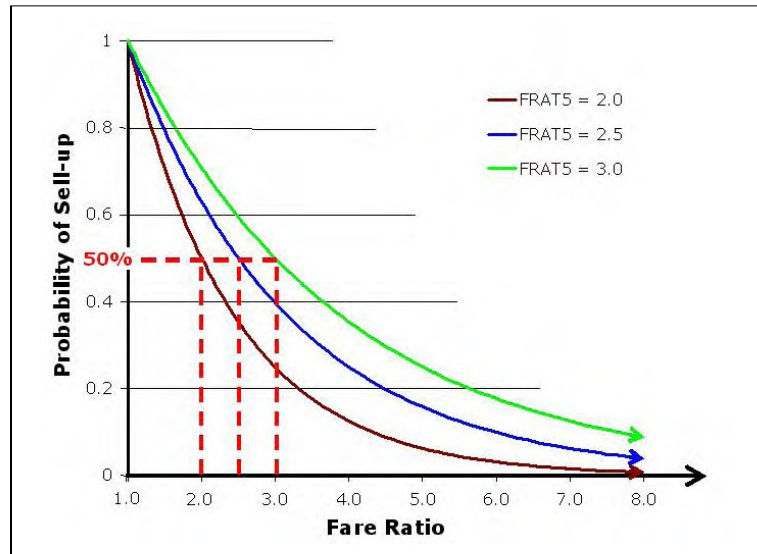
Figure 4-5: Sell-Up Probability in a Time Frame by Frat5 Value[67]

In Soo's[14] thesis, she examined the effectiveness of Fare Adjustment by inputting arbitrary frat5 curves in order to calculate the sell-up probability by time frame. Although PODS allows this input capability, airlines would most likely not arbitrarily input sell-up values into their RM systems. Instead, an airline would utilize its Historical Bookings Database in order to estimate the sell-up probability at a particular time frame.

The methodology for calculating sell-up was initially introduced by Hopperstad, and Cléaz-Savoyen[66] provides a good description of this method in his thesis. The following sections will cover the two sell-up estimators used in the simulation runs in this thesis. Most of the discussion on the Inverse Cumulative (IC) and Forecast Prediction (FP) estimators is based on work from Hopperstad[74] and Guo[75].

4.2.2 Inverse Cumulative Estimator

The more straight-forward of the two estimators used, Inverse Cumulative (IC) operates under the assumption that passengers will buy-down into the lowest available fare class (price-oriented demand). The estimator is based upon the belief that any passenger who is willing to book at some multiple of the base fare Q is also willing to book at any lower fare level. Figure 4-6 shows the process of calculating the observed sell-up probability.

---

[74] Hopperstad, C. (2007). Methods for Estimating Sell-up: Part II. *AGIFORS Joint Revenue Management and Cargo Study Group Meeting*.

[75] Guo, J. C. (2008). Estimating Sell-Up Potential in Airline Revenue Management Systems. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.
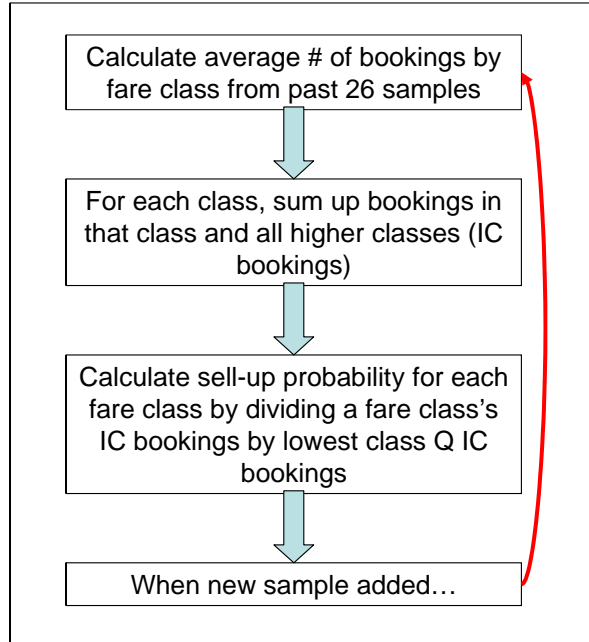
Figure 4-6: Method for Calculating Observed Sell-Up Probability in IC

The first step involves the previous 26 samples as PODS only refers back to the previous 26 departures of a flight for calculations in this thesis. The IC estimator calculates the observed sell-up probability to each fare class by time frame. Therefore, when a new sample is added, that is the number of bookings in the lowest available fare class in the same time frame as all of the other 25 samples being used to generate the sell-up probabilities for a particular time frame in the booking process. The observed sell-up rates are normalized to the 1.0 sell-up rate for the lowest fare class Q.

As Hopperstad[74] and Guo[75] both show, two regressions follow that allow for a frat5 to be obtained for each time frame. The first least squares regression is within the time frame and is used to calculate b such that the following equation is minimized:

$$\sum_{\text{fare ratio}} (psup_{\text{obs, fare ratio}} - e^{-b \cdot (\text{fare ratio} - 1)})^2$$

b is the first estimate of the $econ_{tf}$ and fits the data to a negative exponential. The second regression is across all time frames and $b_{tf}'s$ to calculate $icpt$ (intercept) and $slope$ and minimize:

$$\sum_{tf} ([icpt + slope \cdot tf] - b_{tf})^2$$

Once $icpt$ and $slope$ have been found from the two minimization regressions, a frat5 value can be calculated for each time frame.

$$frat5_{tf} = \frac{-\ln(0.5)}{(icpt + slope \cdot tf)} + 1$$

The sell-up probability for each time frame can then be solved for using the equations introduced in section 4.2.1.

A criticism of the IC estimator is that it is biased. Bookings made in the lowest fare class and the highest fare class are taken as equal although bookings can only occur in the highest fare class if there is high demand for a flight and the RM system has closed lower fare classes. The opposite can be said for bookings in the lowest fare class, but the inherent demand differences with different historical bookings is not taken into account.

4.2.3 Forecast Prediction Estimator

Forecast Prediction (FP) uses the historical bookings from the past 26 samples and converts them into the associated Q forecast for the time frame. These Q forecasts are then used to calculate the estimated sell-up probabilities. One argument for the use of FP is that the ratio of actual bookings to forecast Q bookings represents a correction to the estimated sell-up values, and with a large enough number of corrections an accurate estimate of sell-up can be obtained. The process for estimating sell-up from observed bookings is illustrated in Figure 4-7.
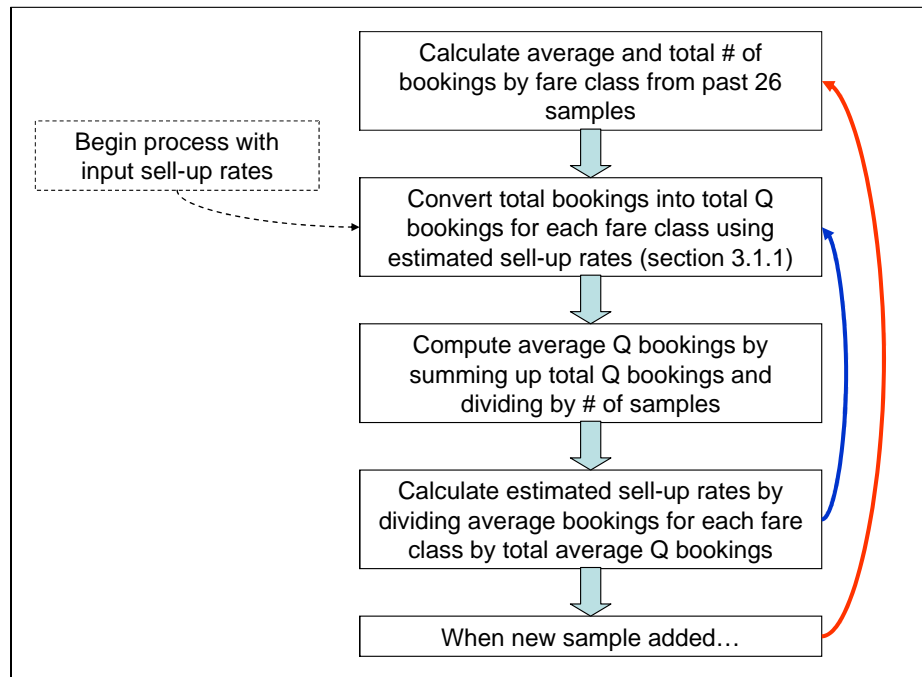

Figure 4-7: Method for Calculating Observed Sell-Up Probability in FP

For the first iteration, sell-up rates inputted by the user are used to calculate total Q bookings for each fare class. The forecast bookings are then used to calculate estimates of sell-up, which are then used when a new sample is added to convert historical bookings into Q bookings. This is where the correction occurs, and a sufficient number of estimated sell-up rate corrections will generate accurate estimates. As with the IC estimator, the observed sell-up rates are normalized to the 1.0 sell-up rate for class Q.

The two regressions performed on the observed estimated sell-up rates are very similar to those in the IC process. The first regression that occurs within a particular time frame adds a scaling factor a and a user-defined weight for each fare ratio. This time, both a and b are selected so as to minimize the equation:

$$\sum_{\text{fare ratio}} wt_{\text{fare ratio}}(psup_{\text{obs, fare ratio}} - a \cdot e^{-b \cdot (\text{fare ratio - 1})})^2$$

However, the second regression across all time frames is identical to IC, with the scaling factor a being dropped from the regression and $icpt$ and $slope$ being solved for by minimizing :

$$\sum_{\text{tf}}([icpt + slope \cdot \text{tf}] - b_{tf})^2$$

Conversion of both $icpt$ and $slope$ into a frat5 is identical to IC, which is then used to calculate the sell-up probability.

4.3   Implementation of New RM Methods in PODS

Chapter 3 described in detail the two new RM methods to be analyzed in this thesis: Hybrid Forecasting (HF) and Fare Adjustment(FA). This section will cover the implementation of these techniques into PODS now that the simulator and related terminology has been introduced.

4.3.1 Hybrid Forecasting

Section 3.1 provided the methodology of Q-forecasting and the incorporation of price- and product-oriented demand into a single forecast. The implementation into PODS by Hopperstad uses the same equations as previously shown, re-calculated at the start of every time frame. Therefore, in Q-forecasting, the sell-up rates calculated by the estimators are used both in the conversion to Q-equivalent bookings and the computation of the mean forecasted demand for a fare class. Q-forecasting operates on price-oriented demand, while pick-up forecasting is performed on product-oriented demand, and they are summed together to create a single forecast for each fare class by time frame.

Currently in PODS, a passenger is considered price-oriented if he booked in the lowest available fare class. Conversely, a passenger is considered product-oriented if the next lower class is available on the same path. Reyes'[67] thesis looked at different "rules" for product demand classification, and the "path rule" outperformed the other options.

The largest problem with HF is identifying which passengers are in fact price-oriented and who are product-oriented. Even for airlines which do not use HF (and may not use the price- and product-oriented terminology), the classification of leisure and business passengers is still a difficult task. The current implementation in PODS is an estimate, as it is obvious that some passengers who book in the lowest available class are specifically looking for that fare product, irrespective of the availability of other fare classes. It is also evident that some passengers who book in a fare class that is higher than the lowest available are not fully product-oriented. For instance, in circumstances where the RM system leaves the lowest fare classes open, this passenger would purchase the lowest available fare class. However, if the RM system has shut down the lower classes and only allows bookings in a few high classes, this passenger may decide he is paying more than he would like to pay anyway (but still below his max WTP), so he will book in a refundable fare class that is higher than the lowest available. Under this scenario, the utility he derives from the refundable fare is greater than the disutility of the higher price. Therefore, this passenger would be considered price-oriented in some cases as he would purchase in the lowest available class, and product-oriented in others as he would choose to book in a fare class that is higher than the lowest available.

This scheme is just one way of classifying demand. Airlines face this same dilemma as the only information used for forecasting demand is booking data. The only way to get a true account of passengers' preferences and characteristics would be through a series of rigorous interviews, which no airline thus far has been willing to do.

4.3.2 Fare Adjustment

The Marginal Revenue Transformation was covered in detail in Section 3.2, where the corresponding equations for the adjusted fare in different fare structures were presented. This section will discuss the implementation of Fare Adjustment (FA) into PODS and the calculation of the PE Cost.

4.3.2.1   FA Formulations in PODS

In PODS, there are two different FA formulations: a continuous marginal revenue formulation (MR), and a discrete formulation (KI, for Karl Isler). The following is the MR formulation implemented in PODS and is the formulation that is used for the simulation runs in this thesis.

$$f_i^{'} = f_i - \frac{f_Q (FA\ FRAT5 - 1)}{-\ln(0.5)}$$

The KI discrete formulation for fare adjustment is shown below.

$$f_i^{'} = \frac{psup_i \cdot f_i - psup_{i-1} \cdot f_{i-1}}{psup_i - psup_{i-1}}$$

4.3.2.2   FA FRAT5

In section 3.2, the actual value input into DAVN's LP with FA was not simply the OD Fare, but the Marginal Revenue, which is the OD Fare – PE Cost. As the MR formulation in PODS shows, the PE Cost is calculated using the fare of the lowest class Q and a FA FRAT5. Since the PE Cost is trying to account for buy-down, as passenger WTP increases, the PE Cost should increase as well, leading to lower classes being closed earlier.

As Cléaz-Savoyen[66] explains in his thesis, the frat5 values used by FA should be lower than the frat5 values used in Q-forecasting. Instead of using two separate frat5's and thus assuming two different WTP's, the two frat5's are linked using a scaling factor.

$$FA\ FRAT5 = 1 + f5scl(FRAT5 - 1)$$

The frat5 used in Q-forecasting is scaled by the value $f5scl$, between 0 and 1, in order to create the FA FRAT5. Any airline in PODS using Q-forecasting (or HF) and FA must decide both on which frat5 to use for forecasting (or sell-up estimator) and which scaling factor to implement for FA that would best represent its passengers' WTP. Figure 4-8 shows a FRAT5C curve and the resulting FA FRAT5 curves with various levels of scaling.
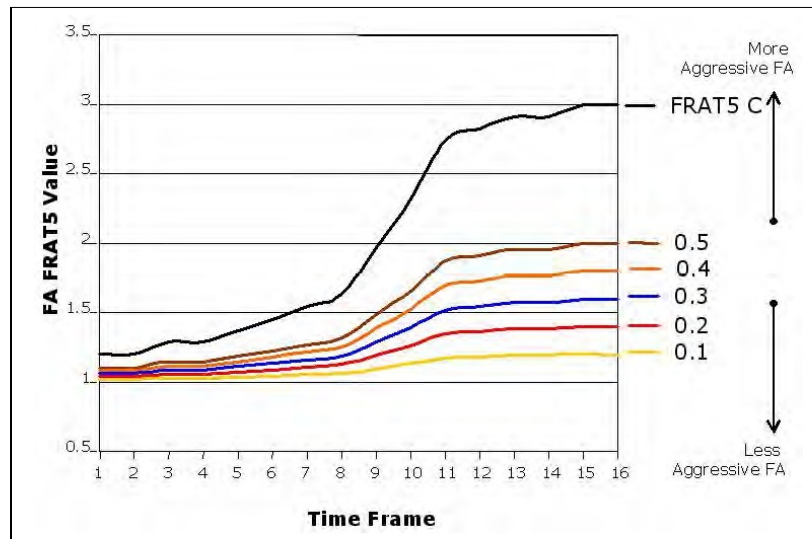
Figure 4-8: PODS FA FRAT5 Curves at Different Scaling Factors

Figure 4-8 shows that scaling the FA FRAT5 downward causes it to become less aggressive. This corresponds to Marginal Revenue values being fed into the LP that are higher than they would have been with a more aggressive scaling factor, therefore increasing the probability that the fare class will remain open. As the results in Chapter 5 show, this scaling factor is used to verify the adjusted fare equation from section 3.2.4.3 that states the more price-oriented demand, the more aggressive the adjusted fare should be. Therefore, in different fare structure environments with differing demand segmentation abilities, different scaling factors (closer to 1.0 for low segmentation fare structures and scaled down further for more restricted structures) will generate the highest revenue.

4.4   PODS Simulation Environments

The environment used in this thesis to evaluate the effectiveness of HF and FA is Network S. This is a four airline (AL1, AL2, AL3, AL4), asymmetric, competitive market meant to model the US domestic market. The network size and markets served differs by airline, while all four airlines operate hubs in the center of the US. AL3 is modeled as a LCC, and the other three airlines are modeled as Network Legacy Carriers (NLC).

Competition is created in the network not only through multiple airlines serving the same market either through direct flights or through a hub, but each airline also serves all 3 competitor hubs in the network. Figure 4-9 to Figure 4-12 depict the route structure of each airline in Network S.
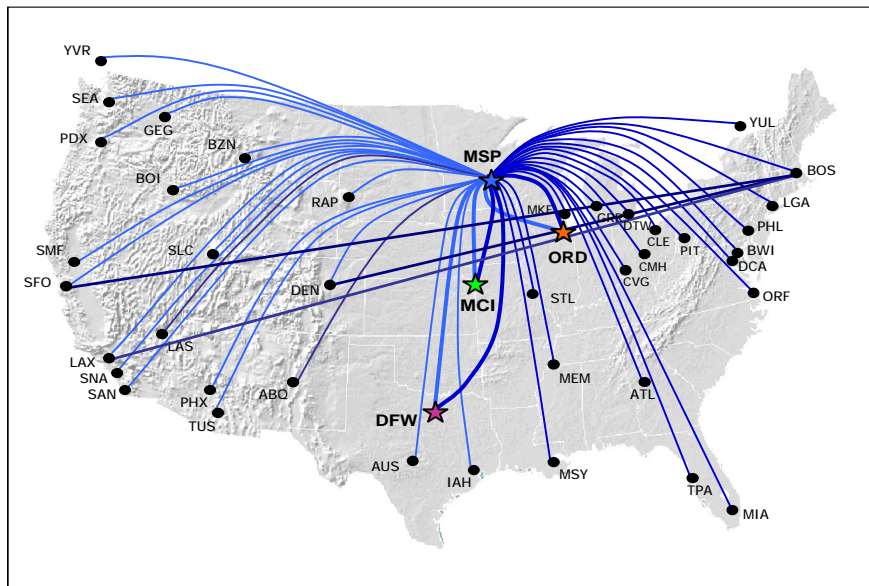
Figure 4-9: Route Structure of AL1
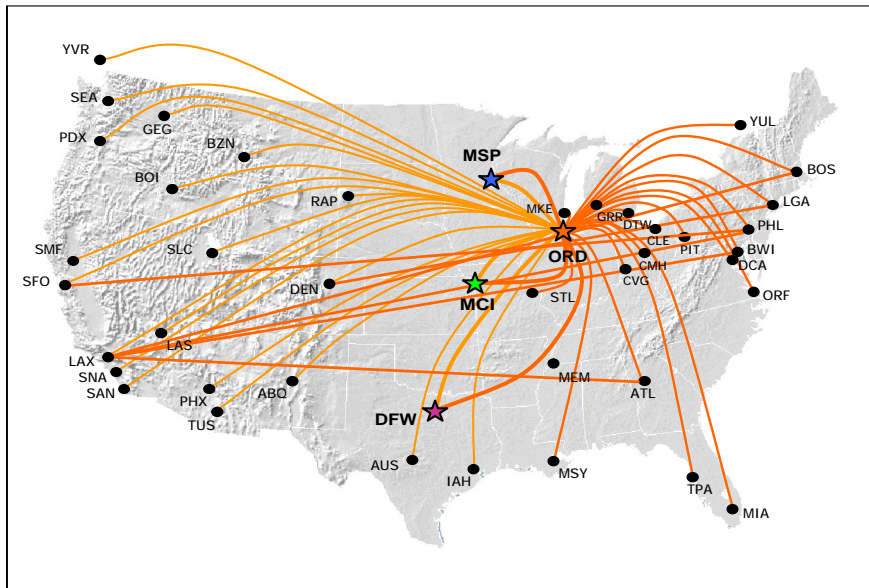
Figure 4-10: Route Structure of AL2
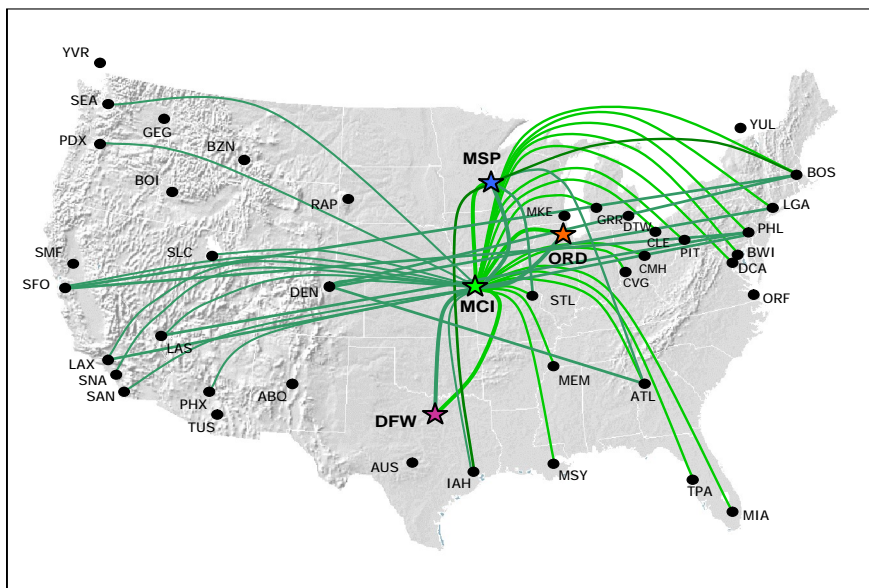
Figure 4-11: Route Structure of AL3
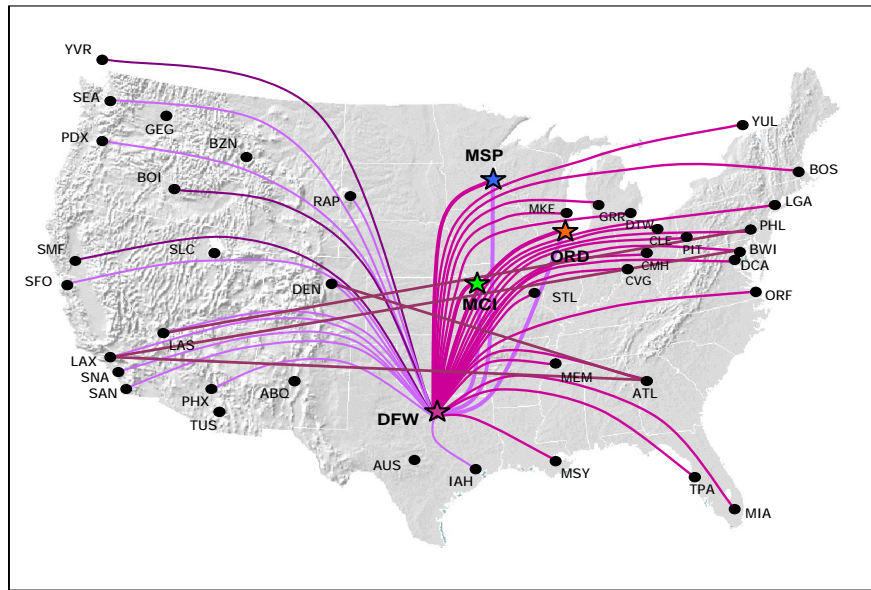
Figure 4-12: Route Structure of AL4

The largest airline, AL1, is based in Minneapolis-St. Paul and serves every market in the network through its hub, as well as three direct flights that bypass its hub. AL2, with its hub in Chicago, is nearly as large as AL1, serving most of the markets in the network and incorporating six direct flights in its network. AL3 is the LCC airline with a hub in Kansas City. It is the smallest airline in terms of markets served (only about half of AL1's markets), but in being modeled as an LCC offers much more point-to-point service than the other three airlines. AL4 is the third legacy airline in the network and is based in Dallas-Fort Worth. It serves the smallest number of markets out of the legacy carriers in Network S.

The route structures for each airline are outlined in Table 4-3.

Airline	# of Origin Cities (Includes Other Hubs)	# of Destination Cities (Includes Other Hubs)	# of Hub Bypass Flights	O-D Markets Served (Local/Connect)
AL1 (MSP)	24	24	3	572 (49/523)
AL2 (ORD)	24	23	6	548 (51/497)
AL3 (MCI)	15	20	19	296 (44/252)
AL4 (DFW)	18	24	4	428 (44/384)

Table 4-3: Network S Route Structure by Airline

Network S is a multiple fare structure environment. Since AL3 is modeled as an LCC, its fare structure is more simplified and fares are compressed. Since AL3 serves 296 markets in direct competition with the other three airlines, the NLC's in Network S match AL3's fare structure in those 296 markets. Therefore, the NLC's operate a more restricted fare structure in the markets where they are not competing against LCC, and match AL3's fare structure in the remaining markets.

In this thesis, two variations of Network S will be used: Network S1 and Network S4. The route structure and schedule of each airline remain the same in both variants. The only difference between the two environments is the set of restrictions on the fare structure employed by AL3 (and matched by the other three airlines).

4.4.1 Network S1

Network S1 is the more realistic of the two networks, given 2008 fare structures in US domestic markets, with less stringent advance purchase and cancellation fee restrictions on the LCC fare structure than the more restricted structure. Also, the LCC structure does not have minimum stay requirements for any fare class. Table 4-4 outlines the fare structure in markets without LCC competition, and Table 4-5 shows the decrease in fares and restrictions in markets with an LCC presence.

Fare Class	Avg Fares	Advance Purchase	Min Stay	Cancellation Fee	Non-Refundable
1	$674.96	None	None	None	No
2	$530.33	3 days	None	Yes	No
3	$385.69	7 days	None	Yes	Yes
4	$257.13	10 days	Yes	Yes	Yes
5	$208.92	14 days	Yes	Yes	Yes
6	$160.71	14 days	Yes	Yes	Yes

Table 4-4: Fare Structure for Markets without LCC Presence in Network S1

Fare Class	Avg Fares	Advance Purchase	Min Stay	Cancellation Fee	Non-Refundable
1	$324.14	None	None	None	No
2	$250.95	None	None	Yes	No
3	$188.21	7 days	None	None	Yes
4	$146.38	7 days	None	Yes	Yes
5	$125.47	14 days	None	Yes	Yes
6	$104.56	14 days	None	Yes	Yes

Table 4-5: Fare Structure for Markets with LCC Presence in Network S1

4.4.2 Network S4

Network S4 is used not as a representation of fare structures airlines will necessarily encounter, but as a "proof of concept" environment in which to test both HF and FA. It is the most extreme case for an LCC fare structure, one in which all restrictions are removed and all passengers will book in the lowest available class. Therefore, HF and FA are tested in this extreme multiple fare structure scenario to evaluate the revenue gains they can produce.

As previously mentioned, the non-LCC fare structure remains identical to Network S1. Only the LCC fare structure is modified, with all restrictions and advance purchase requirements being removed. Table 4-6 shows the new fare structure for these 296 markets.

Fare Class	Avg Fares	Advance Purchase	Min Stay	Cancellation Fee	Non-Refundable
1	$324.14	None	None	None	No
2	$250.95	None	None	None	No
3	$188.21	None	None	None	None
4	$146.38	None	None	None	None
5	$125.47	None	None	None	None
6	$104.56	None	None	None	None

Table 4-6: Fare Structure for Markets with LCC Presence in Network S4

4.5 Chapter Summary

This chapter began by introducing the PODS simulator and the two major components: the Passenger Choice Model and the Revenue Management System. The three Seat Allocation Optimizers that will be used in this thesis were also discussed. Then, the incorporation of sell-up in PODS was covered, as well as covering two sell-up estimators: Inverse Cumulative and Forecast Prediction. The implementation of both Hybrid Forecasting and Fare Adjustment in PODS was discussed, and finally, the Network S1 and S4 simulation environments were described.

The following chapter will provide results of using HF alone and in combination with FA in both network environments to assess their impacts. A new FA implementation will be introduced and tested to evaluate the claim that it reduces the need for scaling to generate the FA FRAT5.

CHAPTER 5
PODS SIMULATION RESULTS

This chapter presents the results of using Hybrid Forecasting (HF) with Fare Adjustment (FA) in a competitive network environment of four airlines. The first half of the chapter is dedicated to Network S1, and the second half contains the same simulation runs in the less restricted Network S4 as well as additional runs when AL2 and AL4 also employ HF and FA. Section 5.1.1 illustrates the revenue gains by incorporating HF alone, and section 5.1.2 discusses the additional gains FA can generate when paired with HF. In section 5.1.3, an alternative FA formulation is described that was developed for an airline which does not scale its FA. The two different network sections (5.1 and 5.2) will conclude with a comparison of four seat allocation optimizers with FA – EMSRb, DAVN, HBP, and ProBP[66].

In order for revenue changes to be accurately attributed to a certain RM technique, all other variables in the simulation must remain constant. AL1 (MSP) uses both EMSRb and DAVN seat allocation optimizers to test HF and FA in these simulation runs, while the other three airlines will maintain the same RM systems: AL2 with DAVN (standard forecasting), AL3 with AT90, and AL4 with DAVN (standard forecasting). Only in section 5.2.2.1.a and 5.2.2.1.b will AL2 and AL4 be given more advanced RM techniques to test the effect on AL1 with HF and FA.

5.1 Network S1

In this section, the effect of HF alone on revenue is examined, as well as the revenue change from HF alone with the addition of FA to HF in Network S1. HF and FA will be tested with the standard Demand Multiplier of 1.00.

In order for different RM techniques to be evaluated against one another easily, they need to be compared to a baseline simulation. Since leg-based forecasting is still prevalent in the airline industry, the baseline simulation run has AL1 using EMSRb with leg-based forecasting (all other airlines using RM methods as described above). Figure 5-1 shows the revenue, load factor and yield of each airline under this "base case" environment. Because of its network and more advanced RM system, AL2 has a slight edge in revenue over AL1. AL3's lack of a large network presence and advanced RM system is evident in its low revenue and yield numbers.
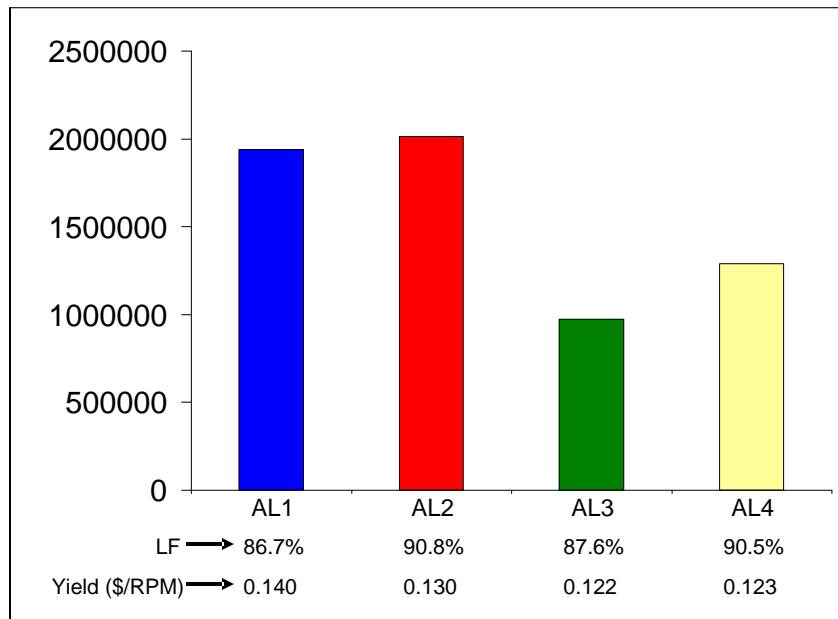
Figure 5-1: Results of Base Case in Network S1

In order to fully ascertain what revenue gains are attributable to Hybrid Forecasting and Fare Adjustment, revenue gains from incorporating path-based forecasting must first be found. Since the HF and FA results in this chapter will be shown while using EMSRb and DAVN, the incremental revenue gains from the base case of incorporating path forecasting with EMSRb and the DAVN optimizer are necessary.
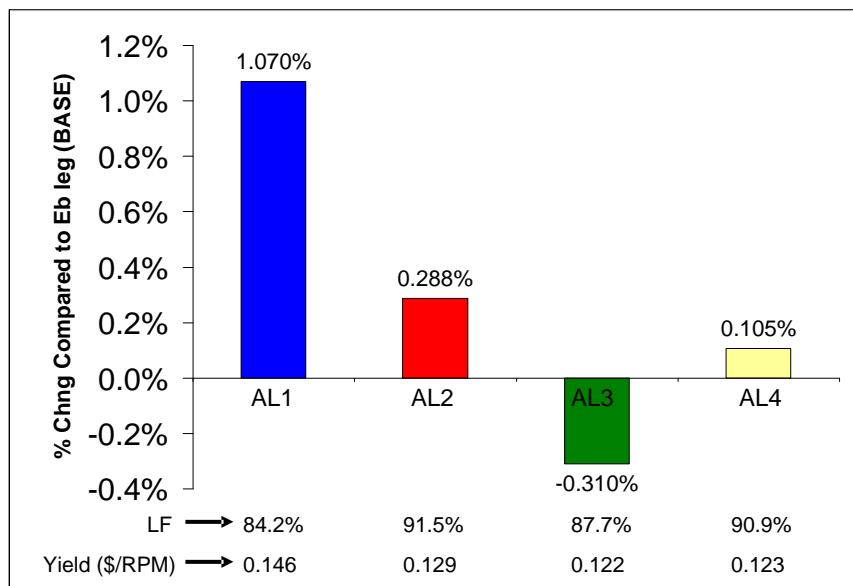
Figure 5-2: Revenue Change with AL1 EMSRb with Path Forecasting

The change from leg to path-based forecasting for AL1 improved its revenue over a percent from the base case (Figure 5-2). This is accomplished through the protection of more seats for late arriving, higher paying passengers, thus the increase in yield and the decrease in load factor. Although the other three airlines kept the same RM system, AL2 and AL4 benefited from AL1's increased aggressiveness by capturing spilled passengers from AL1 and increasing their load factors. AL3, with its simplistic Adaptive Threshold RM system, had a revenue loss of 0.31% from the base case.
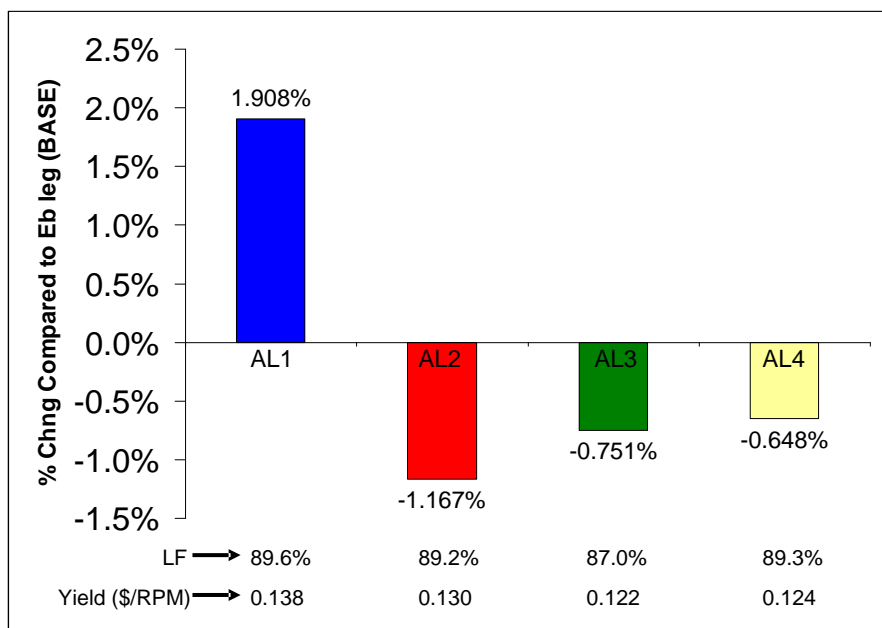
Figure 5-3: Revenue Change with AL1 DAVN

Figure 5-3 illustrates the large revenue gains that can be achieved when switching from a Fare Class Yield Management (section 4.1.2.3.c) optimizer to an O-D optimizer (section 4.1.2.3.d). AL1 sees a revenue jump of slightly less than 2% while the other three airlines lose revenue relative to the baseline simulation. DAVN allows AL1 to increase its load factor by 2.9% and maintain a very similar yield. This load factor increase comes at the expense of the other three airlines, which all see decreased load factors and almost no increase in yield to offset the lower passenger numbers.

5.1.1 Hybrid Forecasting Alone

Now that the revenue gains of optimizers using path-based forecasting have been presented, Hybrid Forecasting can be added to AL1. Recall that HF seeks to classify demand as either price- or product-oriented and thus force the price-oriented demand to sell-up to a higher class than they would have originally booked by closing down lower fare classes sooner.
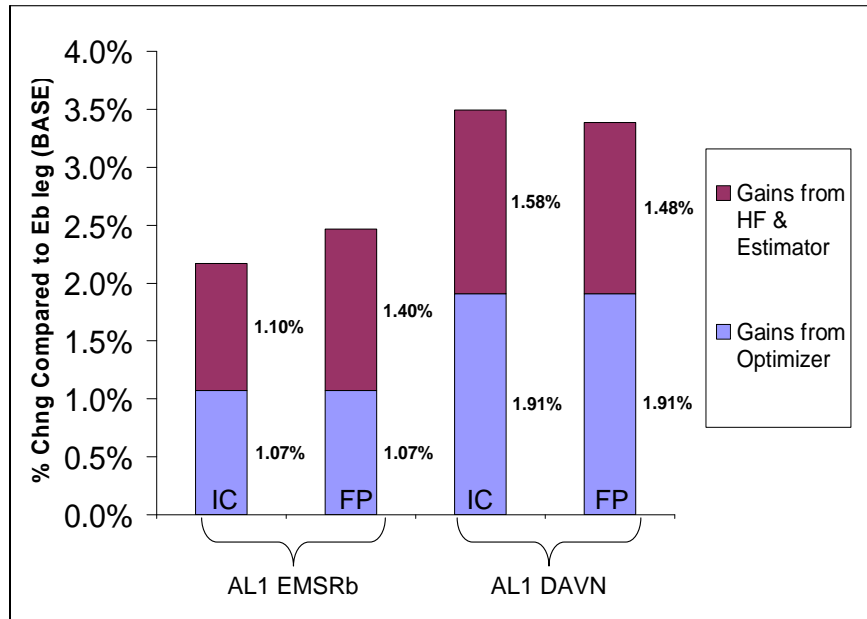
Figure 5-4: Revenue Change of AL1 with HF

As Figure 5-4 depicts, AL1 greatly benefits by the incorporation of HF into its RM system. On top of the revenue gains of path forecasting, HF generates additional 1.1%-1.4% gains for EMSRb and 1.5%-1.6% gains for DAVN, depending on the sell-up estimator used. The DAVN O-D optimizer with HF extends AL1's revenue to almost 3.5% above the base case, which is a very substantial increase for the airline industry which operates under small profit margins. The two sell-up estimators - Inverse Cumulative (IC) and Forecast Prediction (FP) – give relatively similar results when used with HF. A recurring theme in this thesis is first presented here, namely that one estimator is not clearly more effective than the other and results are mixed depending on the other RM techniques being utilized by the airline.
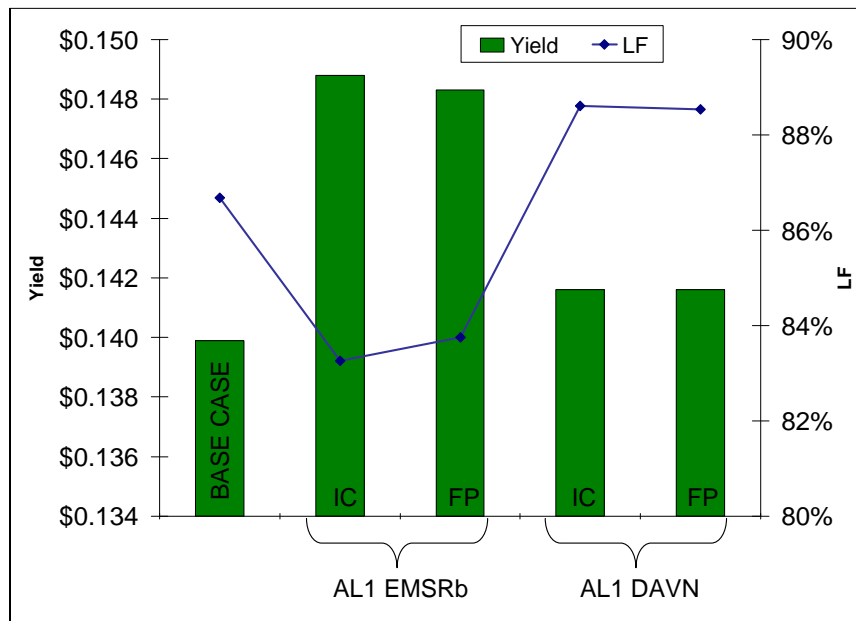
Figure 5-5: Yield and Load Factor with HF Compared to Base Case

With HF forcing more of the price-oriented demand to sell-up, yields for both EMSRb and DAVN increase while load factors decrease slightly. Although HF makes the RM system more aggressive in terms of closing down lower classes, DAVN still maintains a very high average network load factor of just below 89%.

5.1.2 Hybrid Forecasting with Fare Adjustment

In the previous section, DAVN was managing the two fare structures in the network as a coupled structure without the ability for more independent control. This section discusses the additional benefit of FA and the de-coupling of the fare structures. EMSRb and DAVN are evaluated separately as FA was originally developed for use with DAVN and its virtual buckets. The simulation runs with Fare Adjustment were conducted at four scaling levels: 1.0 (no scaling of FA), 0.75, 0.50, and 0.25 (0.25 being the least aggressive, section 4.3.2.2.).

5.1.2.1   EMSRb Path Forecasting

Although FA was developed for implementation with DAVN, it can also be applied to EMSRb when using path-based forecasting. These simulations of FA with EMSRb serve two purposes: first, since many airlines use EMSRb as their Seat Allocation Optimizer, this shows the revenue impacts FA would have on these airlines. Second, it is a way of comparing the intended combination of DAVN with FA to another Seat Allocation Optimizer for which FA was not explicitly designed.

Figure 5-6: Revenue Change of AL1 EMSRb Path Forecasting with HF and FA

The additional benefit of FA added to HF with EMSRb is very small if positive at all, as Figure 5-6 depicts. The IC estimator performs markedly better than FP at all scaling factors for FA, as opposed to HF alo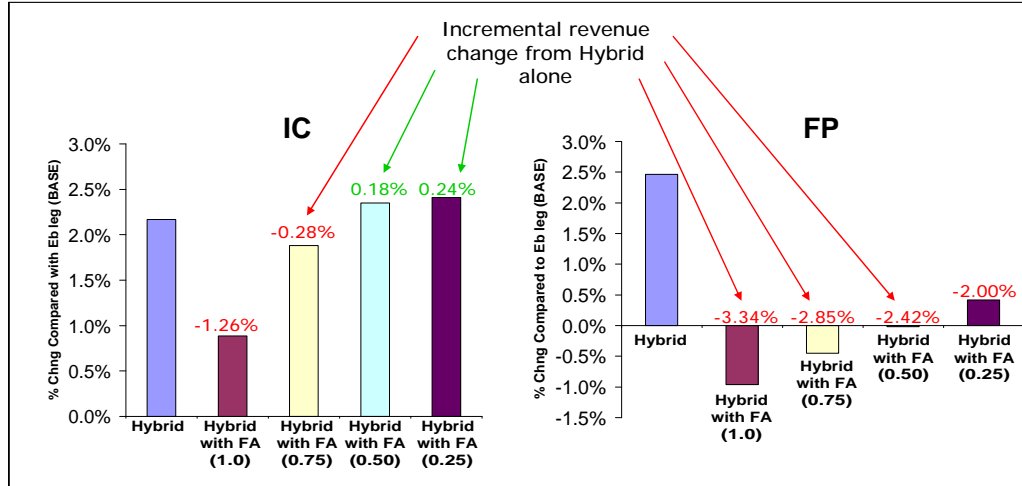ne, where the FP estimator and EMSRb does slightly better than IC. At the lowest FA scaling factors with IC, incremental revenue gains over HF alone are seen, but only on the order of 0.1%-0.2%, whereas the incremental gains of adding HF to an optimizer was above 1% (Figure 5-4). Therefore, in the case of an airline using EMSRb (in a network such as Network S1), the largest jump in revenue will be seen with the adoption of HF, and FA will only give positive results if both an accurate sell-up estimator and FA scaling factor are utilized.

The most interesting observation from Figure 5-6 is the general pattern of revenue, namely that as the FA scaling factors get less aggressive, the revenues increase. This pattern is a product of the fare restrictions in the fare structures of Network S1. Recall from section 3.2.4.3 the equation:

$$f_k^{'} = x f_k + (1-x) f'_k$$

where x is the ratio of product demand to total demand for fare class k. With a higher proportion of product-oriented demand, the weight placed on the adjusted fare for the price-oriented demand is lower, and therefore the overall adjusted fare for the fare class is not lowered significantly from the actual fare price. In order to have a large percentage of product-oriented demand, a fare structure must be able to segment demand relatively well. Network S1 is able to segment demand fairly well and keep business passengers from buying down into the lowest classes, although the more restricted fare structure obviously performs better at this task than the less restricted, LCC fare structure. Since Network S1 is rather restricted, a larger proportion of demand is classified as product-oriented. Therefore, the results from Figure 5-6 support the adjusted fare equation from the Marginal Revenue Transformation, specifically that with a more restricted fare

75

structure, less aggressive Fare Adjustment is needed. This is to be contrasted with the FA results from Network S4 where the LCC fare structure is fully unrestricted.
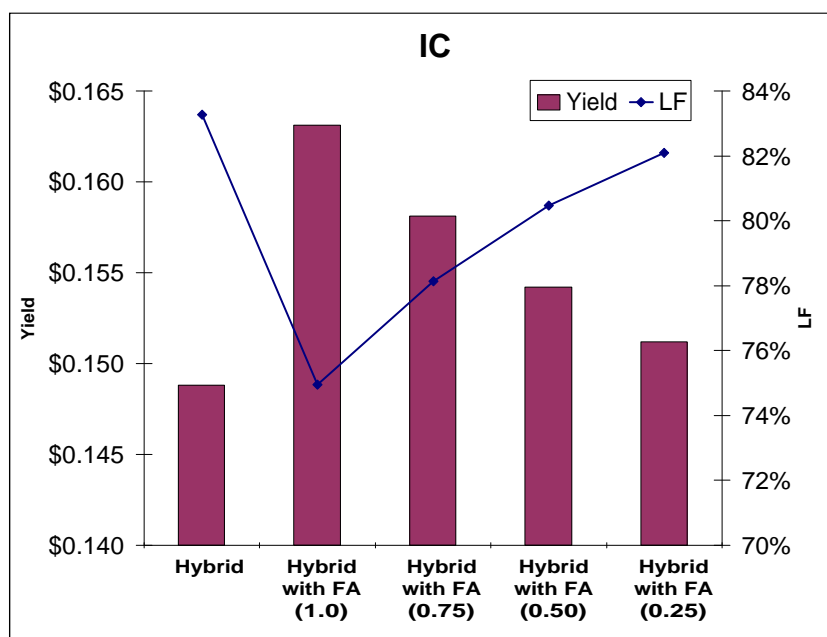
Figure 5-7: Yield and Load Factor of AL1 EMSRb Path Forecasting and IC Estimator with HF and FA

Figure 5-7 illustrates the effects of overaggressive FA scaling in a restricted network. At full Fare Adjustment (scaling = 1.0), the load factor for AL1 drops almost 9%. The yield increase to counteract the passenger loss and the resulting revenue loss from HF is only 0.66%. As the FA scaling factor becomes less aggressive, the load factor continues to rise until it is near its original HF alone level, while the yield correspondingly retreats. With the IC estimator as shown in Figure 5-7, the lower FA scaling factors allow enough lower class passengers back into the system (as opposed to 1.0) that revenues rise slightly, whereas the FP estimator depresses load factors enough that they can never regain their pre-FA level (and rising yields are not enough to compensate). This results in the 2% incremental loss of revenue from HF alone seen in Figure 5-6. The FP estimator creates the same pattern as IC and thus was not shown for analysis purposes.

5.1.2.2   DAVN

Fare Adjustment was applied to EMSRb with path forecasting in order to show the effects of this new RM technique on one of the most widely used Seat Allocation Optimizers in the industry. However, some airlines employ O-D optimizers and more specifically, DAVN. Along with the larger revenue gains generated by the more advanced O-D method (Figure 5-4), the incremental revenue gains of FA are also larger with DAVN as the fare fed into the LP is adjusted for the risk of buy-down and subsequently helps calculate more accurate leg displacement costs.
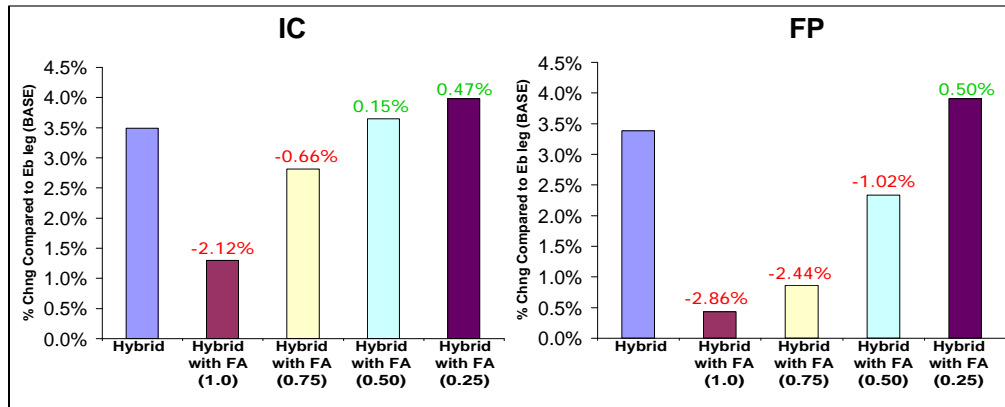
Figure 5-8: Revenue Change of AL1 DAVN with HF and FA

As in Figure 5-6, the pattern of higher revenues with less aggressive FA scaling holds true in Figure 5-8. However, the incremental revenue gains above HF alone for both IC and FP estimators are higher than with EMSRb, as is expected. EMSRb with the IC estimator saw a 0.24% revenue increase over HF alone with FA using a FA scaling factor of 0.25, while this same simulation with DAVN generated a 0.47% increase. The largest difference comes with the FP estimator that performed so poorly with EMSRb and FA, while with DAVN it achieves a 0.50% incremental revenue increase with a FA scaling factor of 0.25.

Even though Fare Adjustment was designed for DAVN and the incremental revenue gains at the best FA scaling factor are 0.50%, the largest revenue jump is again achieved through the implementation of Hybrid Forecasting. Figure 5-4 showed an incremental revenue increase for HF of approximately 1.5% over DAVN with standard forecasting, while Fare Adjustment generates a third of that increase at 0.50% *at the best FA scaling factor*. This is very important as airlines incorporating Fare Adjustment may not know how aggressive to be in order to maximize revenues as is evident with simulation analysis.
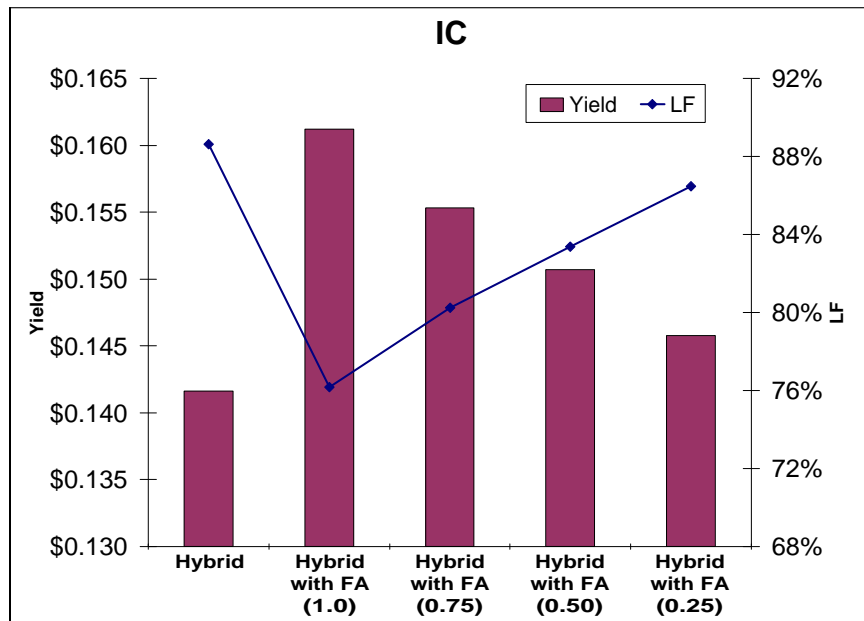
Figure 5-9: Yield and Load Factor of AL1 DAVN and IC Estimator with HF and FA

As in EMSRb, the introduction of full Fare Adjustment precipitates a large drop in the airline's load factor as sell-up is induced through the protection of seats for higher booking classes (Figure 5-9). With the high ratio of product-oriented demand in Network S1, less aggressive Fare Adjustment brings the yield back down with an increase in load factor and generates the 0.47% incremental revenue increase.

5.1.3 Alternate Fare Adjustment Formulation

As the results from section 5.1.2 have shown, incremental revenue gains above HF alone can be achieved with Fare Adjustment and the right FA scaling factor. However, it is still unclear as to how an airline can calculate or estimate the correct FA scaling factor for their respective network in order to maximize revenue. Therefore, an airline may still want to implement Fare Adjustment to combat buy-down while not assuming a FA scaling factor for their network (i.e. using the most aggressive FA scaling factor of 1.0). This alternate FA formulation was developed by Hopperstad[76] to try and generate higher revenues when the airline is not scaling its Fare Adjustment downward.

In the original Fare Adjustment formulation and the one used in all the FA simulation runs up to this point, the final adjusted fare was calculated using the price- and product-oriented demand weighted average of the adjusted fare for the price-oriented demand and

---

[76] Hopperstad, C.. (2008). "Modeling/Programming Update." *PODS Consortium Meeting*, Los Angeles, CA.

the full fare for the product-oriented demand. To make this more evident, recall the final adjusted fare equation from section 3.2.4.3:

$$f_k^{'} = x f_k + (1-x) f'_k$$

This can be rewritten as:

$$f_k^{''} = \frac{prd \cdot f_k + prc \cdot f'_k}{prd + prc}$$

where:
prd is product-oriented demand;
prc is price-oriented demand;
f_k is the full fare;
$f_k^{'}$ is the price-oriented adjusted fare;
$f_k^{''}$ is the final adjusted fare;

In the alternative FA formulation, the price- and product-oriented weights are still used, but instead of appearing in the final adjusted fare equation, they are used to calculate the weighted sell-up to a class.

$$psup_k^{'} = \frac{prc \cdot psup_k}{prc + prd}$$

This weighted sell-up is then run through the PODS conversions in order to incorporate the user-inputted FA scaling factor (for details see Kayser[77]), and finally used directly to calculate the final adjusted fare.

$$f_k^{''} = \frac{psup_k^{'} \cdot f_k - psup_{k-1}^{'} \cdot f_{k-1}}{psup_k^{'} - psup_{k-1}^{'}}$$

where:
$psup_k^{'}$ is the weighted average sell-up to class k;
$psup_{k-1}^{'}$ is the weighted average sell-up to class $k-1$;
f_k is the full fare for class k;
f_{k-1} is the full fare for class $k-1$;
$f_k^{''}$ is the final adjusted fare;

[77] Kayser, M.R. (2008). "Alternative Approach to Hybrid Fare Adjustment." *PODS Consortium Meeting*, Los Angeles, CA.

This new methodology is tested under the conditions for which it was developed, Fare Adjustment with no scaling (1.0).
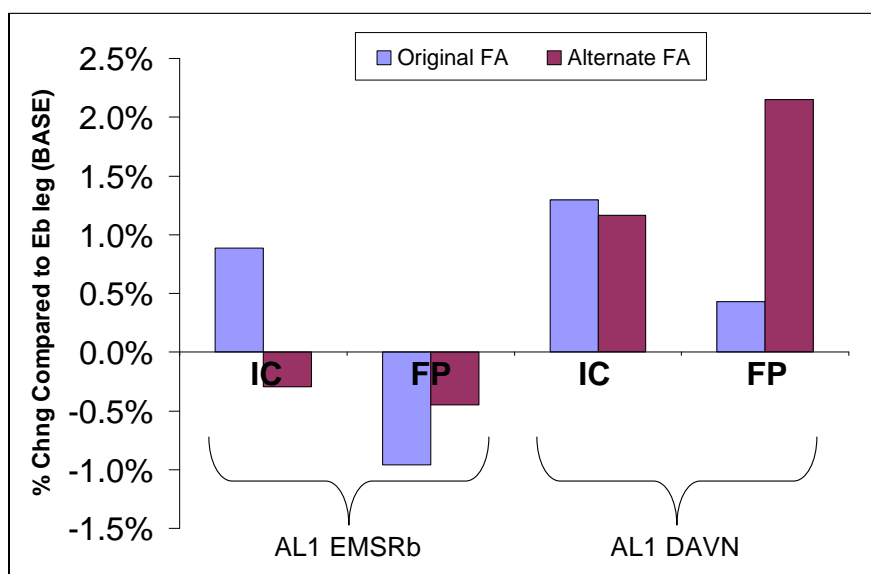

Figure 5-10: Revenue Comparisons of Original FA to Alternative FA Formulation at FA Scaling Factor of 1.0 in Network S1

As Figure 5-10 depicts, the effectiveness of the alternative FA formulation is dependent on the sell-up estimator chosen. The IC estimator performs better with the original formulation, while the FP estimator generates higher revenues when combined with the alternate FA methodology. In fact, the alternative FA method with EMSRb and no scaling gives lower revenues with both estimators than even the simple, EMSRb leg forecasting base case.

However, the combination of the alternate formulation with the FP estimator generates higher revenues than the original FA. The DAVN simulation shows over a 1.5% incremental revenue increase over the original FA method. If an airline which was already using the FP sell-up estimator with HF and wanted to add Fare Adjustment without having to scale, the demand-weighted sell-up FA formulation would generate higher revenues in a network similar to Network S1.

5.1.4 Summary of Best Cases in Network S1

Now that the incremental revenue gains of both Hybrid Forecasting and Fare Adjustment have been presented, this is a good point to step back and compare the best cases for each Seat Allocation Optimizer side-by-side. Heuristic Bid-Price (HBP) and Probabilistic Bid-Price (PBP) are also presented alongside EMSRb and DAVN. The analysis for these two optimizers was left out of this thesis because of their less widespread usage compared to EMSRb and DAVN in the airline industry. Figure 5-11 shows the total revenue gains over the base case broken down by Seat Allocation Optimizer, Hybrid

Forecasting and IC Estimator, and Fare Adjustment. Note that the HBP column does not include Fare Adjustment since this is a "best case" graph and the incorporation of FA with HBP only lowers revenue from HF alone. Also, since the greater number of "best case" simulations involved the IC estimator (including PBP), the results with the FP estimator were omitted.
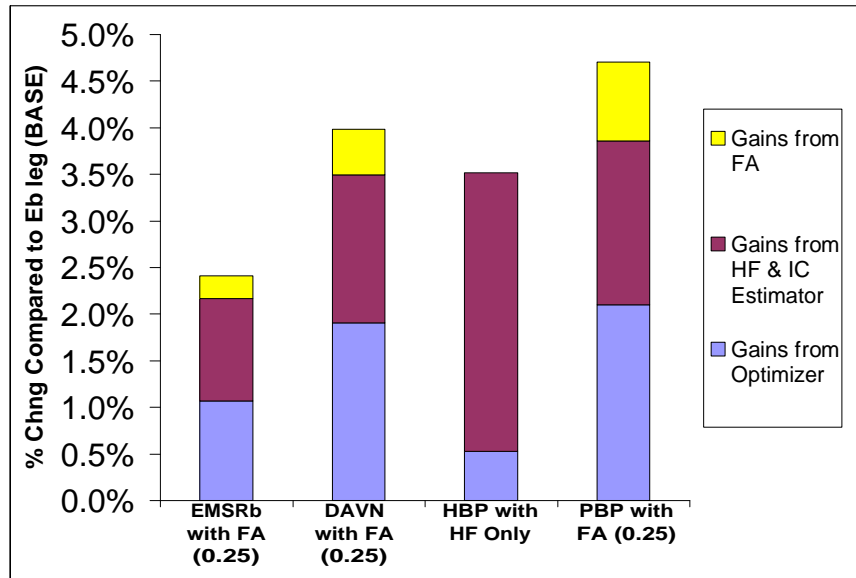
Figure 5-11: Revenue Comparison of Best Cases by Seat Allocation Optimizer in Network S1

Both of the O-D Seat Allocation Optimizers, DAVN and PBP, see the largest revenue gains from Fare Adjustment as well as from the optimizers themselves (Figure 5-11). The combination of HF and FA with DAVN gives a total revenue increase of approximately 4% over the base case, while with EMSRb the increase is almost 2.5%. Although Network S1 has a moderate ability to segment demand and thus reduced need for RM techniques focused on price-oriented demand, revenue gains of approximately 2.5% and 4% for EMSRb and DAVN, respectively, are quite large and attest to the effectiveness of both the forecasting of demand and the ability to manage the different fare structures in the network more independently.
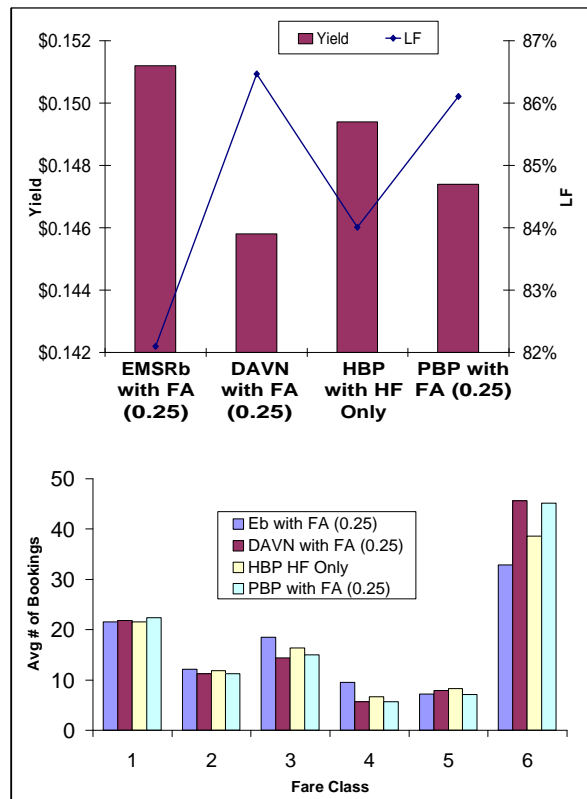
Figure 5-12: Yield, Load Factor, and Fare Class Mix of Best Cases in Network S1

Figure 5-12 presents some counterintuitive results, namely that Fare Adjustment with DAVN, the Seat Allocation Optimizer FA was developed for use with, has the highest load factor even though FA is supposed to shut down lower classes earlier in order to force sell-up. Although it is no surprise that DAVN has the highest number of fare class 6 passengers, it also has a high number of fare class 1 and 2 passengers as well. The lower numbers of bookings come in the middle 3 and 4 fare classes.

The higher revenues coinciding with the higher load factors can partly be attributed to the design of Network S1. High willingness-to-pay, late arriving business passengers will usually book in the higher fare classes because of the fare class restrictions in the network. If the product-oriented demand forecast can be relatively accurate, the number of bookings in the lower classes by price-oriented passengers should be able to be more optimally managed, which may be a reason for the higher load factors and higher revenues for the O-D Seat Allocation Optimizers.

Now that HF and FA have been analyzed in the more restrictive Network S1, they will now be implemented in a network where over half of the markets have no restrictions at all and passengers book solely on price.

5.2 Network S4

Now that the results have been presented for a relatively realistic airline network, HF and FA will be tested in an extreme case for a multiple fare structure environment. The fare structure for markets without an LCC presence remains the same as Network S1, but the markets with LCC competition are completely unrestricted with neither advance purchase requirements nor fare class restrictions. This network was designed to test the effectiveness of RM methods designed to better estimate willingness-to-pay and thus force sell-up, since over half of the markets are unrestricted and fully dependent on the RM system to force bookings in the higher classes. This section contains the same simulation runs as section 5.1, as well as analysis at lower Demand Multipliers (DM) to simulate a weaker demand base and the effects of other airlines using more advanced RM methods.

As in Network S1, the base case for Network S4 is EMSRb leg forecasting for AL1, DAVN with standard forecasting for AL2, AT90 for AL3, and DAVN with standard forecasting for AL4.
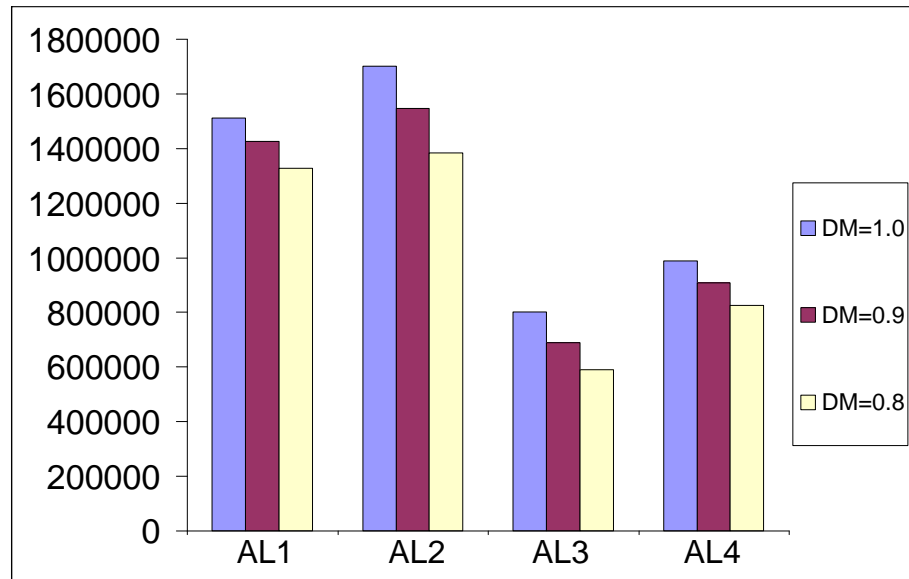
Figure 5-13: Base Case Results in Network S4

Load Factor Yield ($/RPM)	AL1	AL2	AL3	AL4
DM=1.0	91.6%	93.6%	89.1%	93.7%
	0.103	0.106	0.099	0.091
DM=0.9	86.5%	88.4%	85.7%	90.1%
	0.103	0.102	0.089	0.087
DM=0.8	80.5%	80.7%	80.4%	83.6%
	0.103	0.100	0.081	0.086

Table 5-1: Base Case Load Factor and Yield in Network S4

Figure 5-13 and Table 5-1 show the relevant metrics for the base case simulation run at three different demand levels. The difference in the demand segmentation ability of the two networks is evident in the revenue comparison of Figure 5-1 and Figure 5-13. In Network S1, AL1 has total revenue of just under $2M, while in Network S4 with a fully unrestricted fare structure, its revenue drops to approximately $1.5M. This revenue decrease occurs for each airline as they all have markets where they are competing with AL3, the simulated LCC.

Also, as expected, revenues for each airline decrease as the underlying demand is weakened. The proportion of total network revenue for each airline remains the same as demand is lowered, but total revenues are decreased. Lower demand environments are simulated in Network S4 to analyze the revenue improvements of incorporating the higher fare class seat-protecting HF and FA into weaker demand situations. For the rest of the analysis in Network S4, the percent revenue change for each DM is compared to its corresponding DM base case. For example, the light yellow columns in Figure 5-14 (DM=0.8) are percent revenue changes from the DM=0.8 base case.
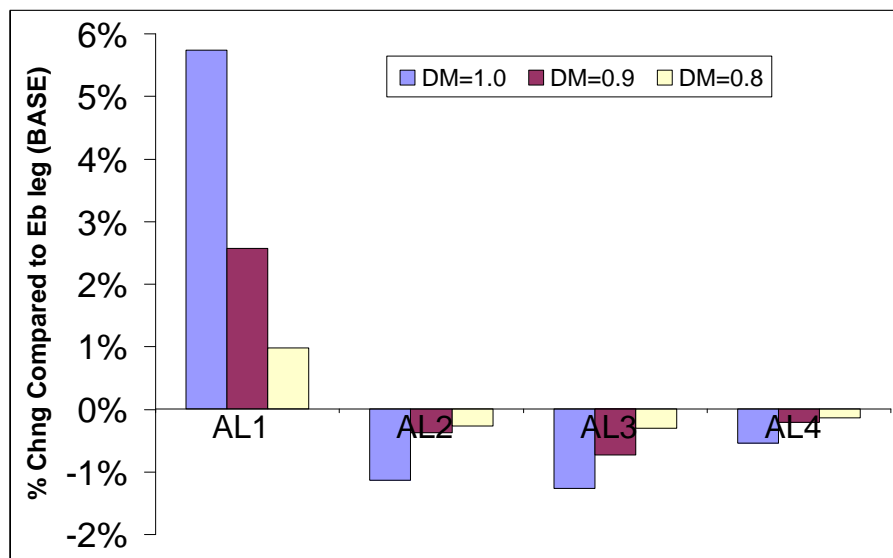
Figure 5-14: Revenue Change with AL1 EMSRb with Path Forecasting

Load Factor Yield ($/RPM)	AL1	AL2	AL3	AL4
DM=1.0	90.6%	94.1%	88.9%	94.0%
	0.110	0.104	0.098	0.091
DM=0.9	86.0%	88.7%	85.6%	90.2%
	0.106	0.101	0.088	0.087
DM=0.8	80.3%	80.8%	80.4%	83.7%
	0.104	0.100	0.081	0.085

Table 5-2: AL1 EMSRb with Path Forecasting Load Factor and Yield

The addition of path forecasting to EMSRb gives a large revenue increase for AL1, almost 6% above leg forecasting. This revenue jump comes at the expense of the other three airlines in the network, which all lose revenue compared to the base case. However, in lower demand environments the effect of path forecasting on AL1 is much less pronounced. In fact, the percent revenue increase is less than 1% at a Demand Multiplier of 0.8. This is intuitive since with less demand, the airline will continue to leave its lower classes open in order to gain more bookings. The addition of path forecasting does not change the RM strategy much at all as forcing sell-up and more accurate forecasts do not play a large role when more bookings are needed at any fare class.

The dwindling effect of path forecasting with lower demands can be seen by comparing Table 5-1 and Table 5-2. At DM=1.0, AL1 sees a 1% drop in load factor and a $0.007 increase in yield due to the effects of EMSRb with path forecasting. However, at DM=0.8, the load factor and yield are almost identical because the demand needed to make path forecasting effective has not materialized. In order to generate the highest number of bookings, AL1 with both leg and path forecasting leave the lower classes open and consequently, a large proportion of the bookings are in fare class 6, the lowest fare class.
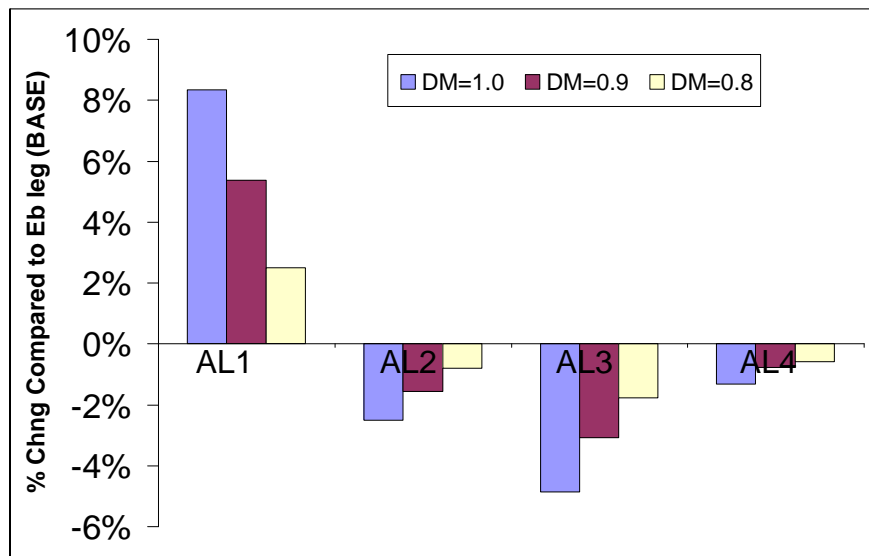

Figure 5-15: Revenue Change with AL1 DAVN

Load Factor Yield ($/RPM)	AL1	AL2	AL3	AL4
DM=1.0	92.9%	92.7%	88.2%	93.3%
	0.110	0.104	0.095	0.091
DM=0.9	88.0%	87.6%	84.6%	89.5%
	0.107	0.101	0.087	0.087
DM=0.8	81.3%	80.5%	79.9%	83.1%
	0.105	0.100	0.080	0.086

Table 5-3: AL1 DAVN Load Factor and Yield

With the DAVN Seat Allocation Optimizer, AL1's revenue increases over 8% from the base case by raising both the load factor and yield (Figure 5-15 and Table 5-3). The added benefit of an O-D optimizer over a FCYM method is evident when comparing load factors and yields: in Table 5-2, EMSRb with path forecasting increased revenues by lowering the load factor but increasing yield. DAVN, on the other hand, is able to generate higher revenues by increasing both the load factor and yield at the same time.

Although all three airlines in the network are adversely affected by AL1's RM improvement, AL3 is the most impacted with a 4.8% decrease in revenue. AL1 not only takes passengers from the other three airlines, but they also take the highest paying passengers as well (lower load factors and yields for other three airlines). As before, the revenue impact of the increased sophistication of the RM system is diminished in lower demand environments. AL1 DAVN at DM=0.8 only achieves a 2.5% increase over the base case while the other three airlines' losses are reduced as well.

5.2.1 Hybrid Forecasting Alone

In this section, Hybrid Forecasting is analyzed in Network S4 with both sell-up estimators. In addition, AL2 and AL4 (both using DAVN) are given Hybrid Forecasting with the IC estimator to simulate the effects on AL1 (with HF) of other airlines in the network using sophisticated RM techniques. Analysis at DM=1.0 is conducted for both the IC and FP estimators, while the lower demand environment simulations are only performed with the IC estimator as revenues with IC are higher for both EMSRb and DAVN at DM=1.0.
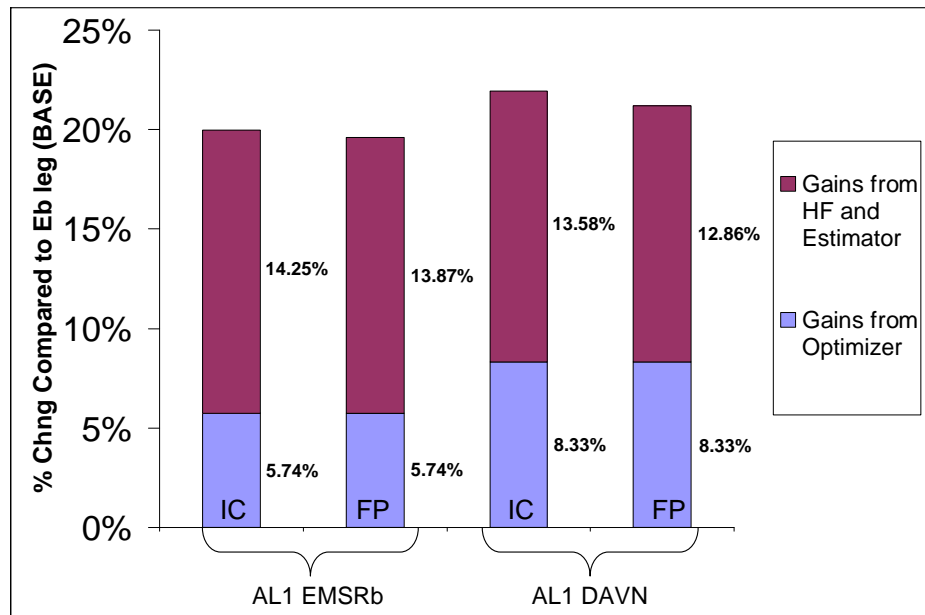
Figure 5-16: Revenue Change of AL1 with HF at DM=1.0
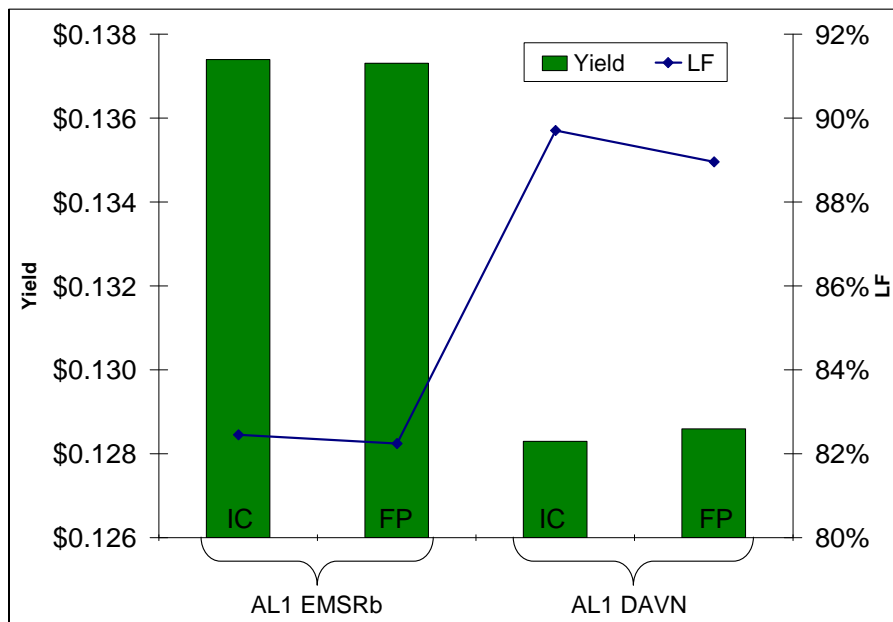
Figure 5-17: Yield and Load Factor of AL1 with HF at DM=1.0

Figures 5-16 and 5-17 give the relevant metrics for the addition of HF to EMSRb and DAVN in Network S4. As compared to Figures 5-4 and 5-5, it is evident that HF provides a much greater incremental benefit in Network S4 when forcing sell-up is even more important. The IC estimator outperforms FP for both optimizers, and the incremental revenue gains of adding HF are all above 12%. These extremely large gains (as compared to Network S1) are due to the fully unrestricted LCC fare structure and the potential revenue increase available with a RM system that can accurately forecast demand and willingness-to-pay in order to force passengers to book in the higher booking

classes. As in Network S1, DAVN has higher revenues than EMSRb with higher load factors and lower yields, although load factors decrease and yields increase for DAVN with the implementation of HF.

Although HF is obviously beneficial with a high level of demand in the network, its benefits may be minimized with weaker demand just as the Seat Allocation Optimizers' benefits were decreased in Figures 5-14 and 5-15.
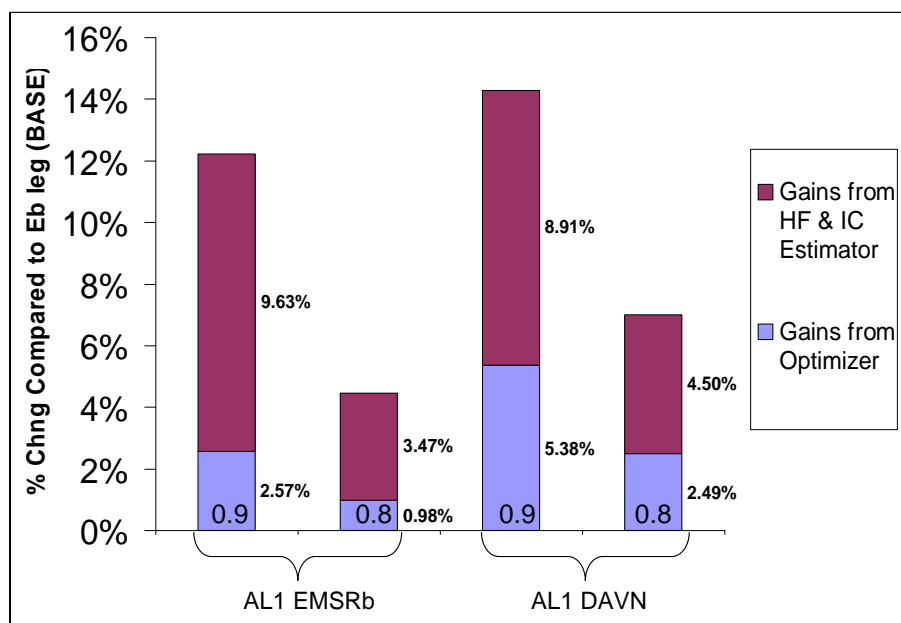
Figure 5-18: Revenue Change of AL1 with HF and IC Estimator at DM=0.9 and DM=0.8

In lower demand environments, the revenue impact of HF is still large, although not quite as great as with DM=1.0. When the Hybrid Forecaster analyzes bookings on previous departures, fewer higher class bookings will be seen and thus a lower forecast for those fare classes. Although the airline wants to have passengers sell-up, the revenue gained from the passengers who will book in higher classes is far outweighed by the revenue lost from passengers who will book elsewhere (or not travel at all) when lower fares are not available.

5.2.1.1   AL2 & AL4 Hybrid Forecasting (IC)

Up to this point, AL2 and AL4 have used an O-D Seat Allocation Optimizer with standard forecasting. Although DAVN performs well with standard forecasting (as evidenced by Figure 5-15), the addition of HF can give substantial revenue benefits (Figure 5-16). In today's low fare environment, airlines are upgrading their RM systems in order to extract higher revenues from passengers while offering the same travel product. Therefore, the current Network S4 simulation environment needs to be updated to reflect the increased sophistication of competing airlines. This is done by allowing

AL2 & AL4 to utilize HF with the DAVN optimizer when AL1 implements HF. Since the IC estimator has been seen to generate the highest revenues in Network S4, all airlines with HF use the IC estimator, and the simulations are run with DM=1.0.
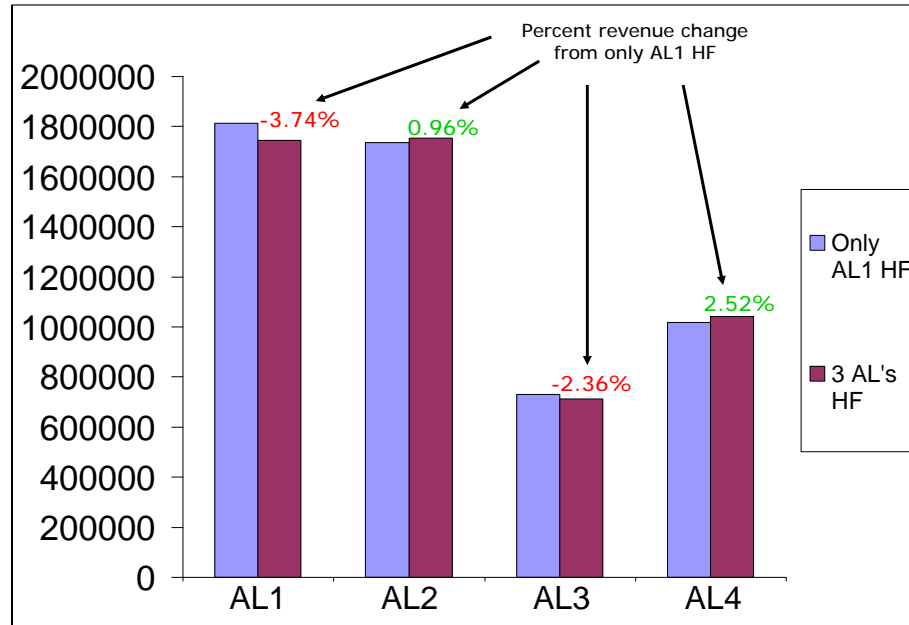
Figure 5-19: Revenue Change with AL2 & AL4 HF at DM=1.0 (AL1 EMSRb)

Load Factor Yield ($/RPM)	AL1	AL2	AL3	AL4
Only AL1 HF	82.5%	96.1%	88.5%	95.6%
	0.137	0.106	0.091	0.092
3 AL's HF	83.8%	94.3%	88.5%	93.7%
	0.130	0.109	0.089	0.096

Table 5-4: Load Factor & Yield Comparison with AL2 & AL4 HF at DM=1.0 (AL1 EMSRb)

With AL2 & AL4 becoming more sophisticated and implementing HF, AL1 loses over 3.5% of its revenue compared to when AL2 & AL4 were only using standard forecasting (Figure 5-19). With the price/product forecasting approach, AL2 & AL4 both see an increase in revenue, with AL4 seeing a greater percentage increase as less of its network is exposed to either AL1 or AL2. AL1's large revenue decrease is largely due to AL2 as it serves almost every market as AL1. AL1's load factor increases as it captures some lower class spill-in passengers (passengers whose first choice was on another airline, but the particular fare class in which they would have booked was closed and they then booked with AL1) from AL2 & AL4 as they shut down their lower classes sooner, but the revenue loss from the lower yield drives down the total revenue (Table 5-4). As shown before, the addition of HF lowers load factors and increases yields as lower classes are closed sooner and passengers are forced into higher fare classes. AL3 is also

negatively affected, losing 2.36% of the revenue it previously enjoyed through a slight decrease in yield.
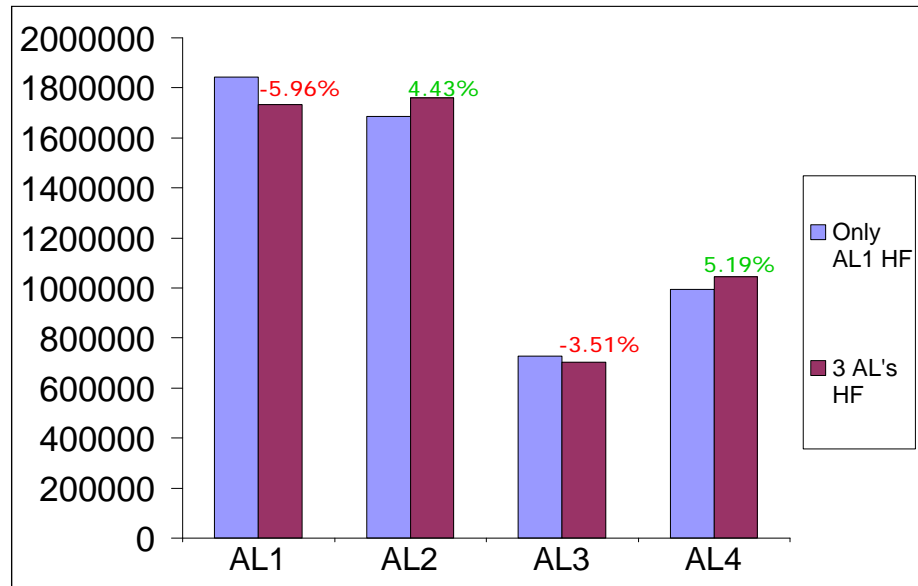
Figure 5-20: Revenue Change with AL2 & AL4 HF at DM=1.0 (AL1 DAVN)

Load Factor Yield ($/RPM)	AL1	AL2	AL3	AL4
Only AL1 HF	89.7%	94.0%	88.1%	94.4%
	0.128	0.105	0.091	0.091
3 AL's HF	90.7%	91.0%	87.9%	92.2%
	0.119	0.113	0.088	0.098

Table 5-5: Load Factor & Yield Comparison with AL2 & AL4 HF at DM=1.0 (AL1 DAVN)

With AL1 DAVN and AL2 & AL4 instituting HF, the percent revenue drop for AL1 is larger than with EMSRb, and likewise the increases for AL2 & AL4 are greater (Figure 5-20). This can be explained by the situation before AL2 & AL4 used HF. When AL1 used EMSRb, it had the highest revenue out of the four airlines, but it was considerably lower than its revenue when using DAVN because of the FCYM optimizer. With only AL1 implementing HF with DAVN, it was able to acquire a larger percentage of the total revenue in the network due to its superior forecasting ability and O-D optimizer. With that being the case, AL1 using DAVN and HF had more to lose in the event another airline increased its RM sophistication. In this scenario, both AL2 & AL4 incorporated HF at the same time and were able to take back some of the network revenue from AL1. Therefore, AL1 drops almost 6% of its revenue from when it was the only airline utilizing HF (Table 5-5). AL2 & AL4 see large increases through better protection for higher-paying, late-arriving passengers (higher yields and lower load factors).

The results in this section support what one may expect with improved forecasting for competing airlines, namely that the airline enjoying the preponderance of revenue will lose some of that revenue, while the airlines improving their forecasters will see an increase in revenue.

5.2.2 Hybrid Forecasting with Fare Adjustment

In Network S1, Fare Adjustment with Hybrid Forecasting led to some incremental revenue gains, but on a much smaller order than the gains from the addition of HF. However, section 5.2.1 showed extremely large gains with HF implemented in Network S4 due to the much lower base case revenues with a fully unrestricted fare structure in the network. Therefore, instituting FA to account for the risk of buy-down and to control the very different fare structures more independently may yield higher revenue gains than previously seen in Network S1.

5.2.2.1   EMSRb Path Forecasting

First, the implementation of FA is analyzed at the standard Demand Multiplier of 1.0 with both sell-up estimators (IC and FP), then lower demand levels are evaluated with the IC estimator only.
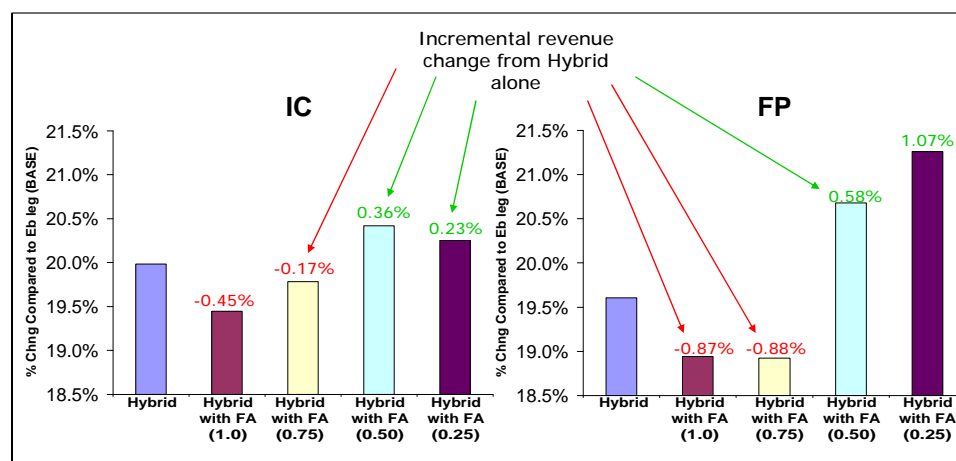
Figure 5-21: Revenue Change of AL1 EMSRb Path Forecasting with HF and FA at DM=1.0

The same conclusion can be made from Figure 5-21 as was made in Network S1, namely that although Fare Adjustment offers some incremental revenue gain, the increase is much smaller than the corresponding incremental increase of the addition of Hybrid Forecasting. Only at a scaling factor of 0.25 using the FP estimator does the incremental benefit eclipse 1%, while the HF gains were all above 12%. However, this is not to say FA offers inconsequential revenue improvement, as tenths of a percent equates to millions of dollars in an industry that deals with yearly revenue in the billions.

With AL1 EMSRb in Network S4, FP actually outperforms FP for the less aggressive FA scaling factors, with 0.25 giving both the highest incremental revenue gain and total revenue of any FA configuration. Also, the revenue pattern exhibited by the two estimators is different, which has not been the case before. One would expect Network S4 to have a higher proportion of passengers classified as price-oriented as all bookings in the LCC fare structure will be in the lowest available class. This will give a greater weight to the adjusted fare for the price-oriented demand in the equation from section 3.2.4.3, ultimately leading to a lower total adjusted fare and a quicker closing of the fare class. The IC estimator reinforces this theory as a more aggressive scaling factor, 0.50, gives the highest incremental revenue increase. However, with FP, the least aggressive scaling factor continues to perform the best, with a 1.07% revenue increase. This discrepancy is more attributable to the estimators themselves and less to the FA methodology.
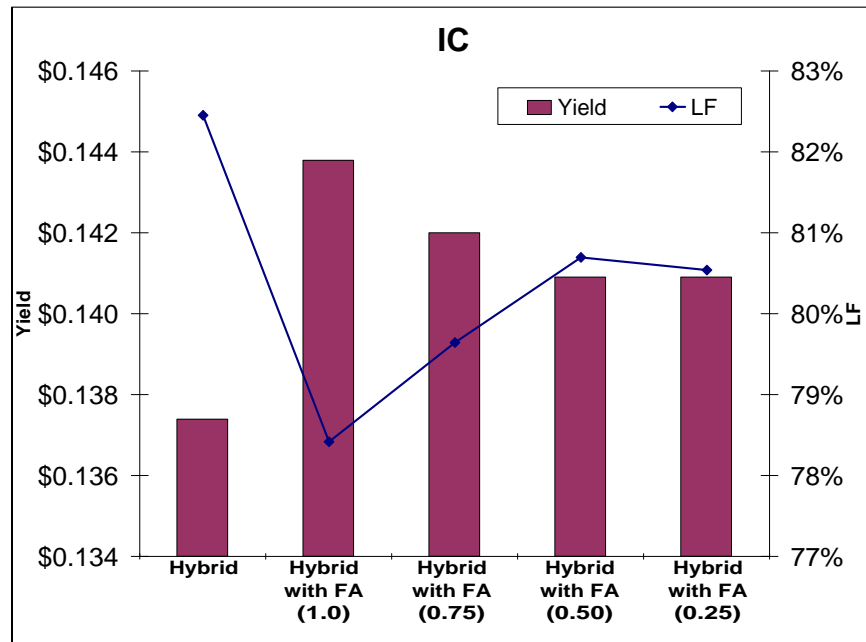
Figure 5-22: Yield and Load Factor of AL1 EMSRb Path Forecasting and IC Estimator with HF and FA at DM=1.0

With the yield and load factor analysis in Figure 5-22, it is easy to determine the reason for the higher revenue at 0.50. As the scaling factor decreases from 1.0 and becomes less aggressive, the load factor rises to just below 81% at 0.50. However, as FA becomes even less aggressive, the load factor does not increase but rather retreats slightly, and the yield stays virtually unchanged, giving 0.50 a higher revenue than 0.25. Although the highest revenue occurring in a higher scaling factor than 0.25 may be expected, one would think that the load factor would continue to increase as the FA becomes less aggressive but a certain yield-load factor combination at a higher scaling factor would generate higher revenues than at the lowest scaling factor. Instead, the load factor tails

off while yield holds relatively constant, making the revenue difference easily identifiable.
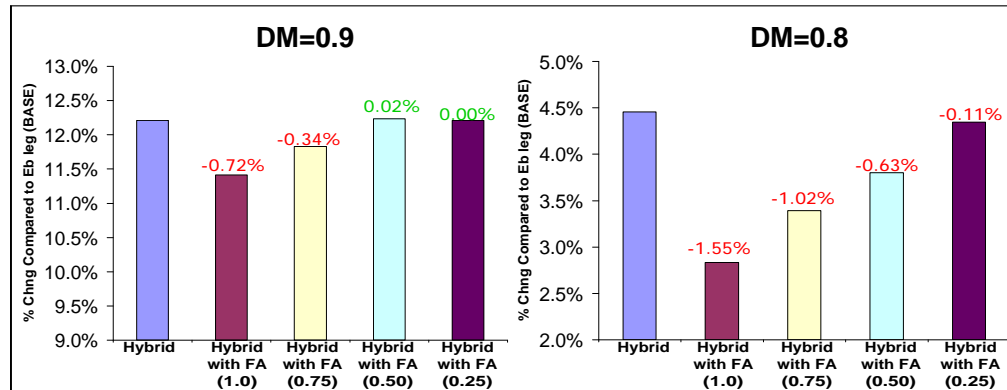
Figure 5-23: Revenue Change of AL1 with HF, FA and IC Estimator at DM=0.9 and DM=0.8

In lower demand environments, the positive effect of FA is virtually negated (Figure 5-23). At DM=0.9, a miniscule 0.02% revenue increase is possible at scaling factor 0.50, but with a weaker demand base at DM=0.8, the aggressive FA methodology does not increase revenues. The revenue pattern from Figure 5-21 continues at DM=0.9, albeit with very small increases, while it breaks down at DM=0.8 and any implementation of FA leads to decreased revenues. It is evident that as the scaling factor converges on 0.0 (HF), revenue will continue to increase until it reaches the HF only level.

5.2.2.1.a   AL1 HF (IC) & FA (0.50), AL2 & AL4 HF (IC) and FA (0.25)

This section, combined with section 5.2.1.2.a, presents the impacts on AL1 when other competing airlines use more sophisticated RM systems. In these analyses, AL2 & AL4 are given Hybrid Forecasting and Fare Adjustment to match what AL1 is utilizing. It is also assumed to be the worst case for AL1 in terms of the other airlines' improvements, namely that since they both are using the DAVN Seat Allocation Optimizer, they implement the best FA scaling factor for use with DAVN (justification for 0.25 scaling factor presented in section 5.2.1.2). Since all three airlines use the IC estimator, AL1 uses a 0.50 scaling factor (per results in Figure 5-21).
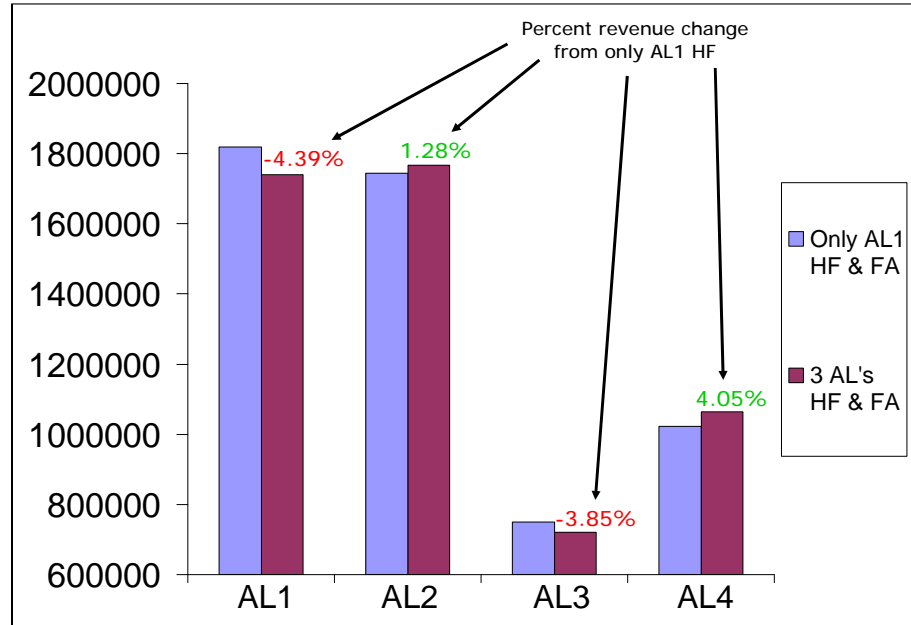
Figure 5-24: Revenue Change with AL2 & AL4 HF and FA at DM=1.0

Load Factor Yield ($/RPM)	AL1	AL2	AL3	AL4
Only AL1 HF & FA	80.7% 0.141	96.1% 0.106	88.7% 0.093	95.5% 0.093
3 AL's HF & FA	82.0% 0.133	94.2% 0.110	88.7% 0.090	92.3% 0.100

Table 5-6: Load Factor & Yield Comparison with AL2 & AL4 HF and FA at DM=1.0

With the addition of HF and FA, AL2 & AL4 protect more seats for higher fare classes and see large revenue gains in the process (Figure 5-25 and Table 5-6). When Figure 5-24 is compared to Figure 5-19, it is evident that most of the revenue gained by AL2 & AL4 is through the incorporation of HF and not FA. FA does stretch the revenue increase a bit, but most of the damage done to AL1 is by the better forecasting methodology. AL3 again is negatively affected as its load factor threshold method is outperformed by more advanced RM techniques.

5.2.2.2 DAVN
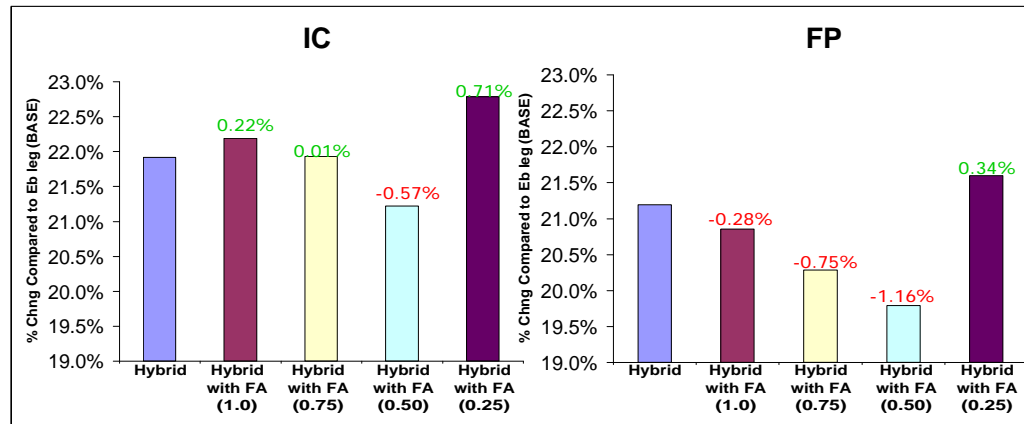
Figure 5-25: Revenue Change of AL1 DAVN with HF and FA at DM=1.0

In Figure 5-25, a new revenue pattern is shown for both estimators when AL1 uses DAVN. As FA becomes less aggressive, revenues decrease through 0.50, and then a large jump is seen at 0.25 that eclipses the HF only revenue and the most aggressive FA revenue. This signifies that the combination of DAVN and HF has already taken into account much of the sell-up behavior of the passengers and has already made those revenue contributions. Although the most aggressive FA does achieve higher revenues with IC, the large jump at 0.25 indicates a slightly more aggressive approach combined with a de-coupling of the fare structures is a more advantageous approach for the airline. Using an O-D Seat Allocation Optimizer with Hybrid Forecasting and an aggressive Fare Adjustment runs the risk of overprotecting seats for higher fare classes and driving the load factor down to levels that higher yields cannot counterbalance.
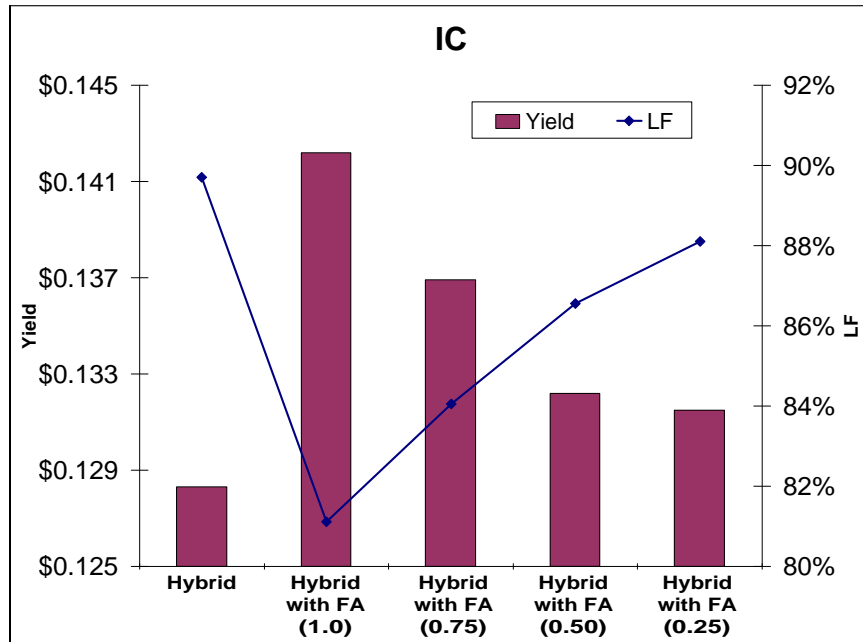
Figure 5-26: Yield and Load Factor of AL1 DAVN and IC Estimator with HF and FA at DM=1.0

As opposed to Figure 5-22, Figure 5-26 is more of the Yield-Load Factor graph that is expected when decreasing the aggressiveness of FA, as the load factor continues to increase and yield continues to decrease. It also provides great insight into the large revenue jump at 0.25. Starting at scaling factor 1.0, the increase in load factor is fairly steady as the scaling factor decreases, with a 1.5% increase between 0.50 and 0.25. However, they decrease in yield isn't as uniform, as the difference between 0.50 and 0.25 is only 0.0007. The revenue jump is a result of this disparity, as the load factor has a large increase while yield stays relatively the same.
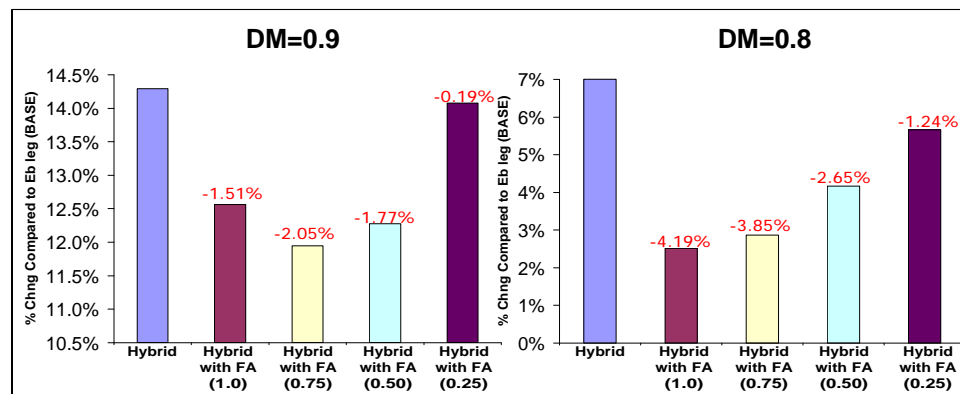
Figure 5-27: Revenue Change of AL1 with HF, FA and IC Estimator at DM=0.9 and DM=0.8

The effect of Fare Adjustment in a lower demand environment is more negatively pronounced in Network S4, as no level of FA aggressiveness produces a better result than HF alone. Weaker demand becomes more of an issue in Network S4 since the LCC markets are unrestricted. Previously in the less-restricted LCC markets in Network S1, even when the airline would have to leave the lower classes open to gain bookings, a certain number of passengers would book in the higher fare classes due to some demand segmentation ability. In Network S4, that ability no longer exists, but trying to shut down classes and force sell-up with a weak underlying demand base will only serve to drive passengers to other airlines and lower revenue.

5.2.2.2.a   AL1, AL2, and AL4 HF (IC) and FA (0.25)

Since AL1 uses DAVN, all three airlines modeled as Network Legacy Carriers will be identical in their RM systems, utilizing Hybrid Forecasting with the IC estimator and Fare Adjustment with a scaling factor of 0.25. These simulation results show the effect on an advanced O-D RM system of competitors' upgrading to a similarly advanced system.

Figure 5-28: Revenue Change with AL2 & AL4 HF and FA at DM=1.0

Load Factor Yield ($/RPM)	AL1	AL2	AL3	AL4
Only AL1 HF & FA	88.1% 0.132	94.2% 0.105	88.4% 0.092	94.5% 0.092
3 AL's HF & FA	89.9% 0.121	90.3% 0.114	88.1% 0.089	91.3% 0.100

Table 5-7: Load Factor & Yield Comparison with AL2 & AL4 HF and FA at DM=1.0

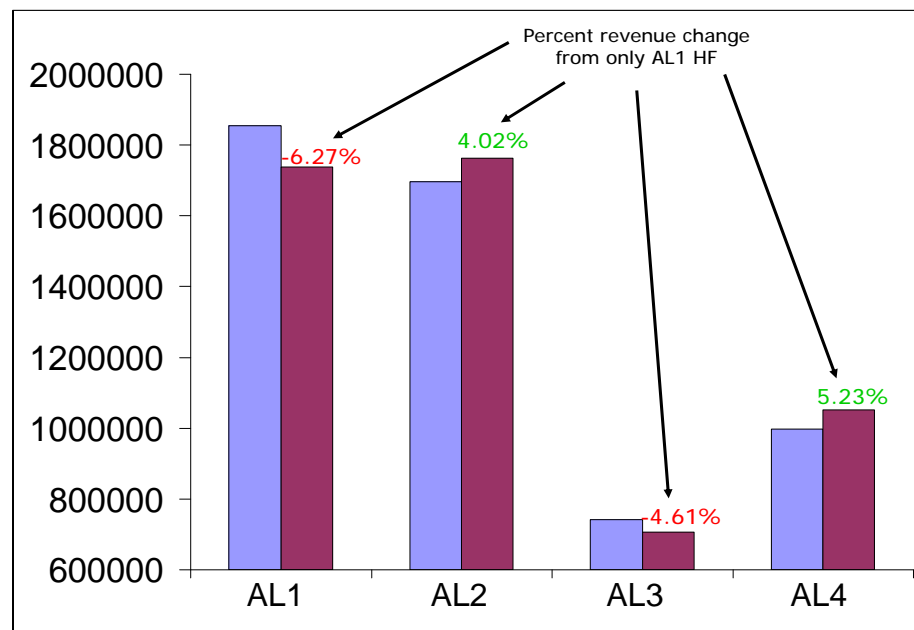As when AL1 used EMSRb and the competition upgraded their RM systems (Figure 5-24), AL1 sees a large drop in revenue even though it is using the exact same RM techniques as its major competition (Figure 5-28). Previously, it was able to extract a high yield by protecting seats for late-arriving passengers and subsequently secure a high revenue. However, the main three airlines are all protecting seats and thus their yields converge toward one another (Table 5-7). AL2 generates the highest revenue out of the four airlines, overtaking AL1 even with a slightly smaller network. The more central location of the ORD hub is one contributor to this revenue advantage.

5.2.3 Alternate Fare Adjustment Formulation

In Network S1, the alternate FA formulation for use without scaling was found to be effective for the FP estimator and not IC (Figure 5-10). With an entire fare structure unrestricted, the weighted sell-up rates should stay reasonably close to the observed sell-up rates as a large proportion of the demand will be classified as price-oriented.

For the analysis of the alternate FA approach, only full-up FA (scaling of 1.0) and the optimal scaling for the given Seat Allocation Optimizer found in this section were run in PODS.
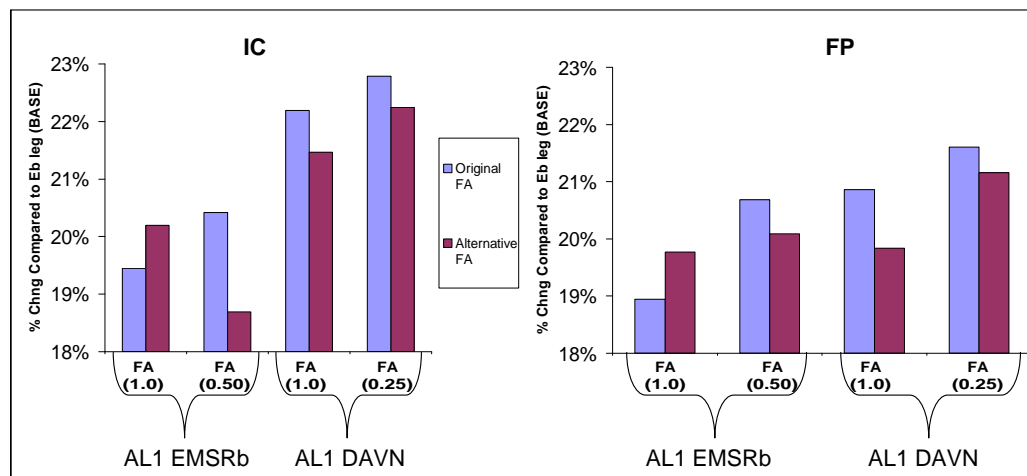

Figure 5-29: Revenue Comparisons of Original FA to Alternative FA Formulation in Network S4

The results of the alternative method in Network S4 are consistent across sell-up estimators. With AL1 EMSRb, the alternative method generates higher revenues when there is no FA scaling occurring, while with DAVN and no scaling, AL1 is better off continuing with the original FA methodology. Figure 5-29 supports the ongoing conclusion that an arbitrary scaling factor can increase revenue in all but one case, where a scaling factor of 0.50 with EMSRb and the IC estimator on the alternative FA approach actually lowers AL1's revenue. In all other cases, however, no matter if the alternative methodology outperforms the weighted adjusted fare approach without scaling or not, higher revenues can be found by scaling the FA approach in use.

5.2.4 Summary of Best Cases in Network S4

In order to get a "big picture" look at the results of different optimizers and RM configurations in Network S4, it is necessary to put them in a side-by-side comparison. As before, these simulation runs return AL2 & AL4 to using DAVN with standard forecasting, as well as the Demand Multiplier to 1.0. HBP and PBP are also added in order for two other Seat Allocation Optimizers to be compared against EMSRb and DAVN. The IC sell-up estimator is used for AL1.
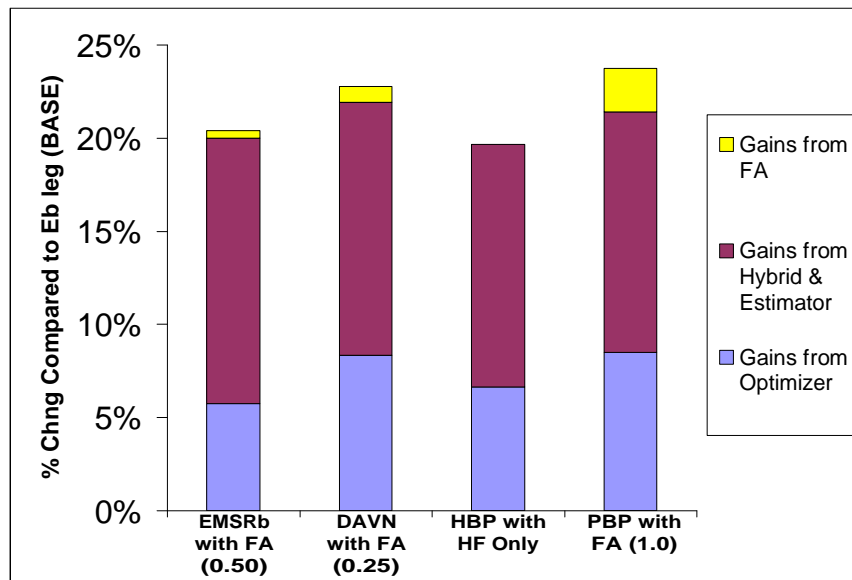
Figure 5-30: Revenue Comparison of Best Cases by Seat Allocation Optimizer in Network S4

As seen earlier and again in Figure 5-30 are the large revenue gains generated by incorporating Hybrid Forecasting. The unrestricted LCC fare structure gives back a lot of revenue from the Network S1 environment, and although the more advanced Seat Allocation Optimizers retrieve over 7% of those losses, HF is where most of the revenue is recouped. Fare Adjustment, especially with the O-D optimizers DAVN and PBP, offers substantial benefit as well with the ability to more independently manage the moderately restrictive fare structure and the completely unrestricted fare structure. Even

more so than in Network S1, the optimal strategy for one fare structure may lead to large revenue drops in markets governed by the other fare structure.

An interesting note is the difference in the FA scaling factors for the two O-D optimizers. Originally, one would expect a more aggressive scaling factor to perform better in Network S4, but throughout the analysis that proved to not be the case, especially with DAVN. However, PBP achieves higher revenues than DAVN with the full-up, most aggressive Fare Adjustment. This is just one of the number of examples which show that not only does the aggressiveness of the FA depend on the fare structures in which it is operating, but also on the other RM techniques being employed by the airline.
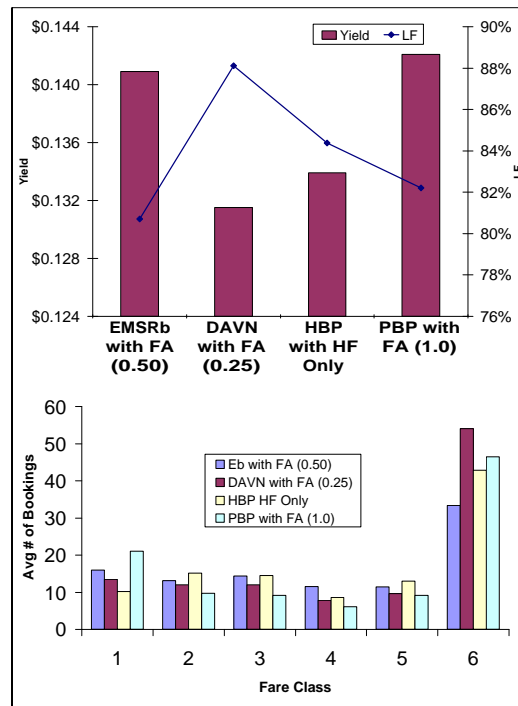
Figure 5-31: Yield, Load Factor, and Fare Class Mix of Best Cases in Network S4

As in Network S1, DAVN generates a high revenue number while having the highest load factor and by far the greatest number of fare class 6 bookings (Figure 5-31). The other O-D optimizer, PBP, achieves an even higher level of revenue, but through a different strategy, by protecting seats for the highest booking class and thereby decreasing the load factor but increasing the yield.

Figure 5-31 can be compared to the metrics presented in Figure 5-12 to evaluate the effect of the unrestricted fare structure. In Figure 5-12, a much larger number of bookings occurred in fare class 1 and 3, while here in Figure 5-31, the first five fare classes are relatively similar in terms of bookings, and a large spike is seen in the lowest booking class. Even though HF and FA are implemented in AL1's RM system, the LCC

fare structure makes it much more difficult to force higher-class bookings, especially in an environment where an LCC is competing on some markets.

5.3   Chapter Summary

In this chapter, the results of simulations focusing on Hybrid Forecasting and Fare Adjustment were presented. These two RM techniques were tested in two different, 4-airline competitive networks, with the restrictions on the markets where an LCC is present being changed. HF and FA were utilized with the EMSRb and DAVN, although final comparisons were conducted with HBP and PBP as well.

In section 5.1, the simulations were run in Network S1, the more realistic of the two networks tested, with the results being summarized in Tables 5-8 and 5-9. With both HF and FA, EMSRb achieved a revenue gain of slightly less than 2.5% over the base case, while DAVN had an increase of over 4%.

	Revenue	% Chng from Base	% Incremental Chng from Previous
Leg Forecasting (Base Case)	1940182	---	---
Path Forecasting	1960941	1.07%	1.07%
Hybrid Forecasting (IC)	1982254	2.17%	1.10%
HF (IC) with Fare Adjustment (0.25)	1986973	2.41%	0.24%

Table 5-8: AL1 EMSRb Revenue Summary in Network S1 at DM=1.0

	Revenue	% Chng from Base	% Incremental Chng from Previous
Standard Forecasting	1977192	1.91%	1.91%
Hybrid Forecasting (IC)	2007906	3.49%	1.58%
HF (IC) with Fare Adjustment (0.25)	2017429	3.98%	0.49%

Table 5-9: AL1 DAVN Revenue Summary in Network S1 at DM=1.0

To test both HF and FA in a more extreme situation, they were given to AL1 in Network S4 with a fully unrestricted LCC fare structure. This is not an altogether realistic environment, but it is useful in evaluating these techniques' ability to account for buy-down and force sell-up in order to maximize revenues. Although the starting base case revenue was lower, HF and FA performed very well as HF provided incremental gains of over 12% and FA of close to 1%. The results are summarized in Tables 5-10 and 5-11.

	Revenue	% Chng from Base	% Incremental Chng from Previous
Leg Forecasting (Base Case)	1510911	---	---
Path Forecasting	1597568	5.74%	5.74%
Hybrid Forecasting (IC)	1812834	19.98%	14.25%
HF (IC) with Fare Adjustment (0.50)	1819419	20.42%	0.44%

Table 5-10: AL1 EMSRb Revenue Summary in Network S4 at DM=1.0

	Revenue	% Chng from Base	% Incremental Chng from Previous
Standard Forecasting	1636821	8.33%	8.33%
Hybrid Forecasting (IC)	1842041	21.92%	13.58%
HF (IC) with Fare Adjustment (0.25)	1855173	22.79%	0.87%

Table 5-11: AL1 DAVN Revenue Summary in Network S4 at DM=1.0

Also in Network S4, HF and FA were tested at lower demand levels. Since these methods are aggressive in trying to force sell-up, the results presented were as expected, namely that as the underlying demand base becomes weaker, the effectiveness of these methods decreases. HF still provided generous revenue improvements (above 3% incremental gains), but FA offered little to no revenue increase, and in most cases it overprotected for higher fare classes and caused the revenue to drop.

Although a FA scaling factor was shown to provide higher revenues than the full, no-scaling FA. However, some airlines (and certainly Revenue Management divisions) do not want to take the risk of assuming a certain scaling factor for their network. For these airlines an alternate FA formulation was tested in order to achieve higher revenues without FA scaling than the original implementation. For the IC estimator, only with EMSRb and in the less-restrictive Network S4 does the alternate formulation outperform the original (Table 5-12).

	Network S1		Network S4	
	Original	**Alternate (% Chng)**	**Original**	**Alternate (% Chng)**
EMSRb with HF (IC), FA (1.0)	1957342	1934456 (-1.76%)	1804729	1816047 (0.63%)
DAVN with HF (IC), FA (1.0)	1965306	1962787 (-0.98%)	1846155	1835253 (-0.59%)

Table 5-12: AL1 Revenue Results with Alternate FA Formulation

Finally, simulations were conducted to evaluate the effects that increased sophistication of competitors' RM system had on AL1. In Network S4, this meant that AL2 and AL4 would be given HF when AL1 was using HF, and HF and FA along with AL1. As expected, AL2 & AL4 captured back some of the revenue that AL1 was able to generate when it was the only airline using these techniques. AL1's losses were 3%-6.5% from its previous revenue totals. Therefore, implementing HF and FA still achieves revenue gains, but the improvement is smaller when other airlines are using similarly sophisticated RM systems.

CHAPTER 6
CONCLUSIONS

Since many Network Legacy Carriers (NLC) now employ multiple fare structures across their networks, this thesis tested the ability of two new Revenue Management techniques, Hybrid Forecasting (HF) and Fare Adjustment (FA), to generate higher revenues in these multiple fare structure environments. In the more realistic Network S1, the combination of HF and FA achieved revenue increases above a standard leg-based forecasting approach of nearly 4% in the best case, while in the "proof-of-concept" Network S4, these methods together generated gains of over 20%.

The beginning of this thesis presented the reader with an overview of airline revenue management as well as some of the traditional revenue management models and the need for new techniques. Many of the conventional methods assume that demand for a class is independent of demand for any other class. Although this assumption is almost never valid, in the years following deregulation the use of this assumption could be justified with the ability to segment demand through advance purchase requirements and booking restrictions in the fare structure. However, the compressed fare ratios and reduction of fare product restrictions brought about by the emergence of LCC's invalidates the demand independence assumption. In fact, traditional RM techniques applied to these simplified fare structures cause a spiral down of not only bookings into the lowest class, but ultimately revenue as the bookings database records more lower class bookings and fewer high fare class seats are protected. Clearly, there is a need for new RM methods for use in this new environment.

Chapter 3 described two RM techniques designed for use with simplified fare structures. Hybrid Forecasting breaks demand into two categories, price- and product-oriented, and forecasts demand separately for each group in a fare class. Fare Adjustment allows an airline with multiple fare structures in its network (normally a more-restricted structure for markets without an LCC presence and matching the LCC on markets where they both compete) to more independently manage the seat inventory between the two fare structures. The objective of this thesis was to evaluate the effectiveness of these two methods in a four-airline, asymmetric, competitive network.

The Passenger Origin-Destination Simulator (PODS) was introduced in Chapter 4, and its Passenger Choice Model and Revenue Management System were discussed. Next the seat allocation optimizers (EMSRb, DAVN, AT90) and sell-up estimators (IC, FP) employed in the simulation runs were covered. Instead of using user-defined inputs to calculate the sell-up probability at different times during the booking process, this thesis utilized two sell-up estimators, Inverse Cumulative and Forecast Prediction, to calculate the probability of sell-up from the Historical Bookings Database. This makes the analysis more realistic as an actual airline would not assume a sell-up value. The conversion of HF and FA from theory development by Fiig et al. in Chapter 3 to implementation in PODS by Hopperstad was also illustrated.

The two simulated networks used in this thesis were identical in route structure for each airline, with the only difference being the fare structure employed on markets with an LCC presence. Network S1 was a more realistic network as the less-restricted fare structure still employed advance purchase requirements along with cancellation and refund restrictions. Network S4 was a "proof of concept" network designed to assess the new RM techniques' effects on an airline in the most extreme case, when a fare structure is completely unrestricted.

Finally, the results of the simulation runs were presented in Chapter 5, focusing on the effects Hybrid Forecasting and Fare Adjustment had on an airline in a multiple fare structure, competitive environment. In an effort to eliminate the need for FA scaling (a method used to vary the aggressiveness of the Fare Adjustment), an alternate FA formulation was introduced and tested in both network environments.

6.1 Summary of Findings

In general, the introduction of HF and FA into a RM system gives significant revenue improvement over the Seat Allocation Optimizer with standard forecasting. The impact of both techniques depends upon the restrictions present in the multiple fare structures as well as the underlying demand in the network. However, Fare Adjustment is sometimes detrimental as it overprotects seats and causes a large drop in load factor. For an airline implementing FA to achieve its highest possible revenue, a FA scaling factor is used to coordinate FA's aggressiveness with the demand segmentation ability of the fare structures.

The first network where HF and FA were applied was Network S1. Because the fare structures segment demand moderately well, the base case (EMSRb leg) gave a high total revenue for AL1. However, the implementation of HF and FA with EMSRb and DAVN increased revenues 2.4% and 4.0%, respectively, over the base case. A summary of the findings in Network S1 is shown in Figures 6-1 and 6-2 (FA scaling factor shown is the best case for that particular Seat Allocation Optimizer).
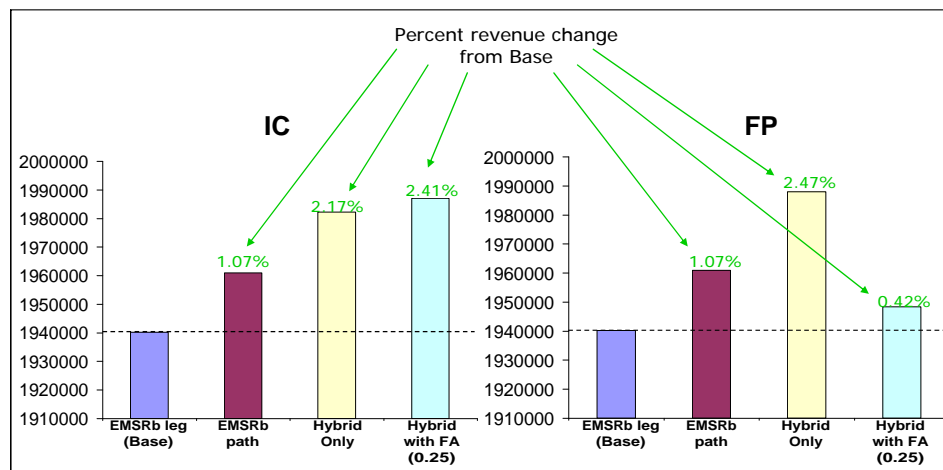

Figure 6-1: Revenue Results for AL1 EMSRb at DM=1.0 in Network S1
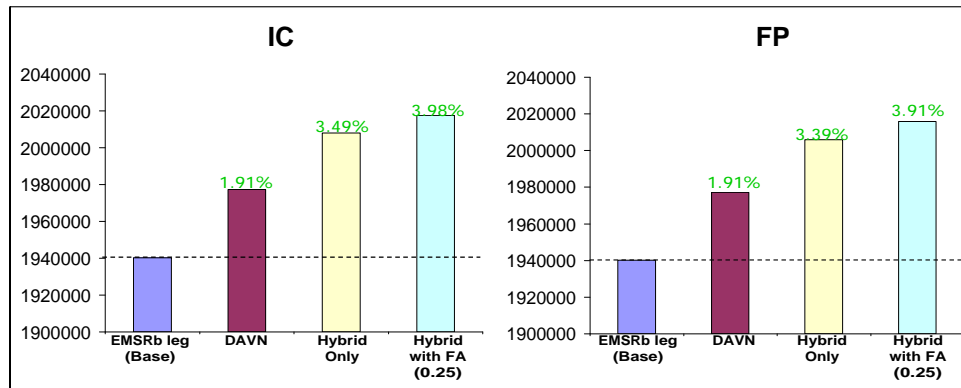
105

Figure 6-2: Revenue Results for AL1 DAVN at DM=1.0 in Network S1

HF and FA give revenue gains over the base case, with DAVN seeing a greater percent increase in revenue along with higher total revenue. This is to be expected, especially with Fare Adjustment considering FA was designed for use with the virtual classes constructed in the DAVN optimizer. With the demand segmentation ability of Network S1, the best FA scaling factor for each optimizer is 0.25, which is the least aggressive FA approach tested. In all but one case in Network S1, the addition of FA increases AL1's revenue, with larger percent increases seen with DAVN. Even though FA has a positive impact, the greatest revenue enhancer (between the two new RM techniques tested) comes from the forecasting approach of Hybrid Forecasting, which uses the price- and product-oriented classification to more accurately determine future bookings.
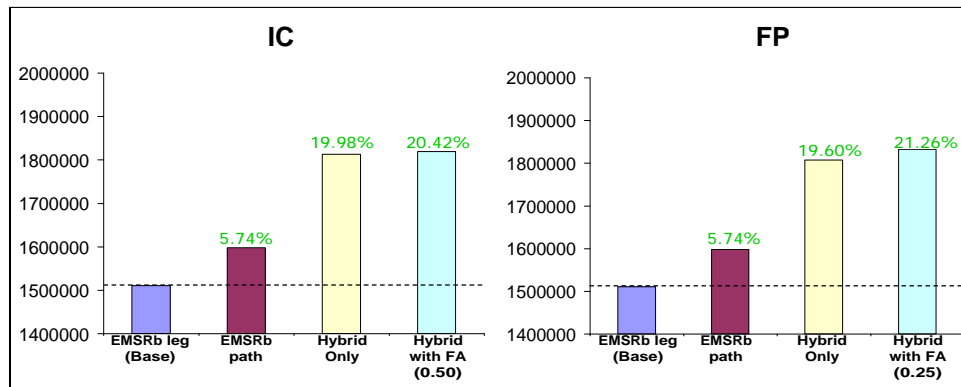

Figure 6-3: Revenue Results for AL1 EMSRb at DM=1.0 in Network S4
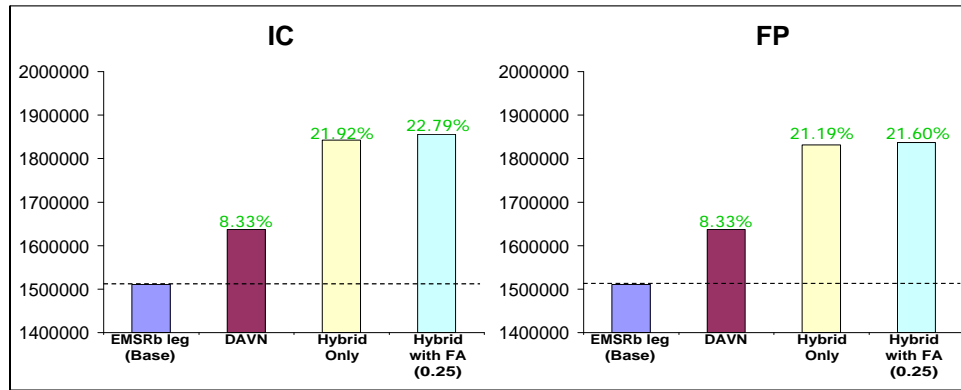
Figure 6-4: Revenue Results for AL1 DAVN at DM=1.0 in Network S4
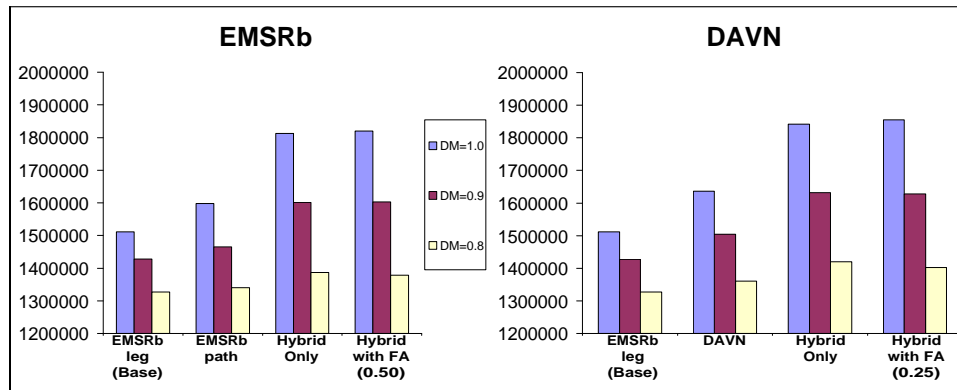
Figure 6-5: Revenue Results for AL1 with IC Estimator at Different Demands in Network S4

The overall pattern of large increases in revenue with the addition of Hybrid Forecasting and smaller increases with Fare Adjustment continues in Network S4, although the percent revenue increases over the base case are on the order of 20% (Figures 6-3 and 6-4). This occurs because of the reduced revenue of the base case due to the fully unrestricted LCC fare structure and the inability of EMSRb leg to force bookings in the higher classes in those markets. The effect of HF and FA in lower demand environments was also tested, and as Figure 6-5 shows, the positive gains HF and FA generate decrease substantially as there is less demand overall in the network.

Figure 6-6: Revenue Results for AL1 when AL2 & AL4 Match AL1 RM Techniques in Network S4

After showing that HF and FA generate large revenue increases for AL1, we wanted to examine the effect on AL1 when AL2 and AL4 (both using DAVN) matched AL1's RM techniques. Therefore, when AL1 implements HF, AL2 and AL4 are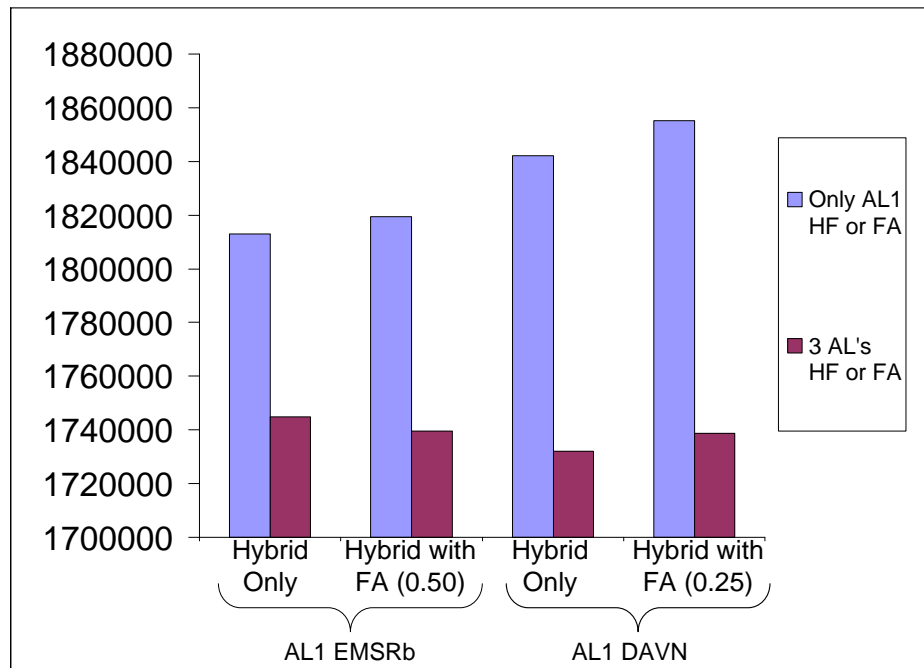 also given HF, and when AL1 uses HF and FA with its optimal FA scaling, AL2 and AL4 are also given HF and FA with their optimal FA scaling (0.25 in Network S4). Much of the revenue that AL1 obtains when it's the only airline utilizing the new RM techniques is given back to AL2 and AL4 when they are given more sophisticated systems. When AL1 is using EMSRb, it loses over 3.5% of its revenue, and when it is using DAVN, it loses over 6% of its revenue. Given the percentage of revenue loss is larger for DAVN, this is largely due to the higher revenue when AL1 is the only airline utilizing HF and FA with DAVN. It secures an even larger amount of revenue with HF and FA while utilizing DAVN and thus is in a position to have more of that revenue redistributed to other airlines in the network.

Finally, an alternate FA formulation was tested in order to eliminate the need for FA scaling. The results of this new methodology without scaling were mixed in both networks. In Network S1, the alternate FA formulation outperformed the original method with both EMSRb and DAVN using the FP estimator, while in Network S4, EMSRb with the new methodology attained higher revenues irregardless of the sell-up estimator. However, in every case, the original FA formulation with an appropriate FA scaling factor generated higher revenues than either methodology without scaling.

The results presented in this thesis allow some general conclusions to be made:

- A lack of restrictions and demand segmentation ability in a network tends to lower revenues, but allows Hybrid Forecasting and Fare Adjustment to have a greater positive effect (as a percentage of baseline revenues)
- The largest revenue gains are obtained through the implementation of Hybrid Forecasting, and smaller gains may be possible with the correct scaling of Fare Adjustment
- The underlying network demand must be strong or the effectiveness of these new RM techniques becomes negated
- When competitors are using sophisticated RM systems, the addition of Hybrid Forecasting and Fare Adjustment still offer revenue improvements, but the gains are smaller than when competitors employ more sophisticated RM systems

6.2 Future Research Directions

In order to make the results presented in this thesis more applicable to the airline industry, sell-up estimators were used instead of user-defined sell-up values. However, throughout the thesis, the results for a simulation run using either IC or FP were continuously mixed in terms of revenue gains. A more rigorous method of estimating a passenger's willingness-to-pay and subsequent sell-up probability is vital to the effectiveness of the HF and FA methods. At this point, it is difficult to determine how much of an effect on revenues the sell-up estimators are having, but it is certainly obvious that for different fare structures and different employed RM techniques, different sell-up estimators give the highest revenues.

Although Hybrid Forecasting was responsible for large jumps in revenue in almost every case shown, there is still an inherent price-oriented demand bias in the forecaster. Currently, any passenger who books in the lowest open class is categorized as price-oriented, even though this passenger may have been looking for this exact fare product and it just happened to be the lowest open. Refinements to this forecasting approach need to be developed which can overcome this bias and categorize passengers more effectively. If this new forecaster can be created, it should be combined with Fare Adjustment and tested in PODS to compare with the current Hybrid Forecasting approach.

In Chapter 3, the Marginal Revenue Transformation was described and the interaction between the price-/product-oriented demand and the adjusted fare was shown. Although the results from the simulation runs generally supported this equation, different Seat Allocation Optimizers would maximize their revenues with different FA scaling factors, or different levels of FA aggressiveness. More testing needs to be done to investigate the relationship between the Seat Allocation Optimizer, fare structures present in the network, and the level of FA aggressiveness. Either this research needs to be conducted, or an effective FA methodology that does not require the use of scaling needs to be developed. However, with the contributions of Fiig et al.[69] and the results from this and other theses, finding the interactions of the network and the airline's own RM system with the aggressiveness of the Fare Adjustment seems much more within reach.

Bibliography

Alstrup, J., S. Boaz, O.B.G. Madsen, R. Vidal, V. Victor. (1986). Booking Policy for Flights with Two Types of Passengers. European Journal of Operations Research, Volume 27, pp. 274 -288.

Barnhart, C., P. P. Belobaba, A. R. Odoni. (2003). Applications of Operations Research in the Air Transport Industry. *Transportation Science*, Volume 37, Issue 4, pp. 368-391.

Belobaba and Hopperstad. (2004). "Q investigations – Algorithms for Unrestricted Fare Classes." *PODS Consortium Meeting*, Amsterdam.

Belobaba, P. (2008). "Overview of PODS Simulation Models and RM Systems." *PODS Consortium Meeting*, Los Angeles, CA.

Belobaba, P. P. (1987). Air Travel Demand and Airline Seat Inventory Management. Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA.

Belobaba, P. P. (1992). The Revenue Enhancement Potential of Airline Revenue Management Systems. *ASTAIR Proc. Adv. Software Tech. Air Transport,* London, U.K.

Belobaba, P. P. (1998). Airline differential pricing for effective yield management. *The Handbook of Airline Marketing*, D. Jenkins (ed.). The Aviation Weekly Group of the McGraw-Hill Companies, New York, NY, pp. 349-361.

Belobaba, P. P. (1998). The evolution of airline yield management: fare class to origin-destination seat inventory control. *The Handbook of Airline Marketing,* D. Jenkins (ed.). The Aviation Weekly Group of the McGraw-Hill Companies, New York, NY, pp. 285-302.

Belobaba, P., C. Hopperstad. (2004). Algorithms for Revenue Management in Unrestricted Fare Markets. Presented at the Meeting of the INFORMS Section on Revenue Management, Massachusetts Institute of Technology, Cambridge, MA.

Belobaba, P., C. Hopperstad. (2004). Algorithms for revenue management in unrestricted fare markets. Presented at the Meeting of the INFORMS Section on Revenue Management, Massachusetts Institute of Technology, Cambridge, MA.

Belobaba, P.P., L.R. Weatherford. (1996). Comparing Decision Rules that Incorporate Customer Diversion in Perishable Asset Revenue Management Situations. *Decision Sciences*, Volume 27, Issue 2, pp. 343.

Bender, A. R., F. J. Stephenson. (1998). Contemporary Issues Affecting the Demand for Business Air Travel in the United States, *Journal of Air Transport Management*, Volume 4, Issue 2, pp. 99-109.

Boeing Airplane Company. (1997). Decision Window Path Preference Methodology Description. Seattle, WA.

Boyd, E. A., I. C. Bilegan. (2003). "Revenue Management and E-Commerce." *Management Science,* Volume 49, Issue 10, pp.1363-1386.

Boyd, E. A., Kallesen, R. (2004). The science of revenue management when passengers purchase the lowest available fare. *Journal of Revenue and Pricing Management*. Volume 3, Issue 2, pp. 171-177.

Bratu, S. J-C. (1998). Network value concept in airline revenue management. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

Brumelle, S. L. and McGill, J. I. (1988). Airline Seat Allocation with Multiple Nested Fare Classes. Paper presented at the Fall ORSA/TIMS Conference, Denver, CO. Also presented at the University of British Columbia, 1987.

Carrier, E. (2003). Modeling Airline Passenger Choice: Passenger Preference for Schedule in the Passenger Origin-Destination Simulator (PODS). Master's Thesis, Massachusetts Institute of Technology, Cambridge, MA.

Cary, D. (2004). Future of Revenue Management: A view from the inside. *Journal of Revenue and Pricing Management,* Volume 3, Issue 2, pp. 200-203.

Clarke, M., B. Smith. (2004). Impact of Operations Research on the Evolution of the Airline Industry. *Journal of Aircraft*. Volume 41, Issue 1, pp. 62-72.

Cléaz-Savoyen, R. L. (2005). Airline Revenue Management Methods for Less Restricted Fare Structures. Master's Thesis, Massachusetts Institute of Technology, Cambridge, MA.

Cooper, W. L., T. Homem-de-Mello, A. J. Kleywegt. (2004). Models of the Spiral-down Effect in Revenue Management. Working Paper, Department of Mechanical Engineering, University of Minnesota, Minneapolis, MN.

Curry, R. E. (1990). Optimal Airline Seat Allocation with Fare Classes Nested by Origin and Destinations. *Transportation Science*. Volume 24, Issue 3, pp. 193–204.

Cusano, A. J. (2003). Airline revenue management under alternative fare structures. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

De Lollis, B. (Aug 22, 2006). Fare Changes Give Fliers Weekends Back. *USA Today*. Retrieved January 28, 2007, from the World Wide Web: http://www.usatoday.com/travel/flights/2006-08-21-saturday-stays-usat_x.htm.

Delta Airlines. (Jan 5, 2005). Delta Slashes Everyday Fares up to 50 Percent as Airline Introduces SimpliFares^{TM} Nationwide. Delta Airlines Press Release. Retrieved January 28, 2007, from the World Wide Web: http://news.delta.com/article9584.html.

Donnelly, S., James, A., Binnion, C. (2004). bmi's response to the changing European airline marketplace. *Journal of Revenue and Pricing Management*. Volume 3, Issue 1, pp. 10-17.

Dunleavy, H., Westermann, D. (2005). Future of airline revenue management. *Journal of Revenue and Pricing Management*. Volume 3, Issue 4, pp. 380-282.

Elliott, C. (Jan 17, 2006). When Fliers Benefit From Airline Bankruptcy. *New York Times*. Retrieved February 19, 2008, from the World Wide Web: http://www.nytimes.com/2006/01/17/business/17soff.html.

Fiig, T., Isler, K. (2004). "SAS O&D low cost project." *PODS Consortium Meeting*, Minneapolis, MN.

Fiig, T., K. Isler, C. Hopperstad, P. Belobaba, (2008). Optimization in Airline Fare Structures with Restricted and Unrestricted Products. Draft Paper. PODS Consortium.

Forsyth, P. (2003). Low-cost Carriers in Australia: Experiences and Impacts. *Journal of Air Transport Management*, Volume 9, Issue 5, pp. 277-284.

General Accounting Office. (1999). Airline Deregulation: Changes in Airfares, Service Quality, and Barriers to Entry. Report to Congressional Requesters. GAO/RCED-99-92. Washington, D.C.

Gittell, J. H. (2003). *The Southwest Airlines way: using the power of relationships to achieve high performance.* McGraw-Hill, New York, NY.

Gorin, T. O. (2000). Airline revenue management: sell-up and forecasting algorithms. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

Gorin, T. O. (2004). Assessing low-fare entry in airline markets: impacts of revenue management and network flows. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA.

Gorin, T., P. Belobaba. (2004). Revenue management performance in a low-fare airline environment: insights from the Passenger Origin-Destination Simulator. *Journal of Revenue and Pricing Management.* Volume 3, Issue 3, pp. 215-236.

Guo, J. C. (2008). Estimating Sell-Up Potential in Airline Revenue Management Systems. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

Hopperstad, C. (2007). Methods for Estimating Sell-up: Part II. *AGIFORS Joint Revenue Management and Cargo Study Group Meeting*.

Hopperstad, C.. (2008). "Modeling/Programming Update." *PODS Consortium Meeting*, Los Angeles, CA.

Isler and Fiig, (2004). SAS O&D Low Cost Project, PODS presentation, Minneapolis-St Paul.

Kayser, M.R. (2008). "Alternative Approach to Hybrid Fare Adjustment." *PODS Consortium Meeting*, Los Angeles, CA.

L'Heureux, E. (1986). A New Twist in Forecasting Short-term Passenger Pickup. *26th AGIFORS Annual Symposium Proceedings*, Bowness-on-Windemere, England, pp. 248–261.

Lee, A. Y. (1998). Investigation of Competitive Impacts of Origin-Destination Control using PODS. Master's Thesis, Massachusetts Institute of Technology, Cambridge, MA.

Lee, S. (2000). Modeling Passenger Disutilities in Airline Revenue Management Simulation. Master's Thesis, Massachusetts of Technology, Cambridge, MA.

Littlewood, K. (1972). Forecasting and Control of Passenger Bookings. *12thAGIFORS Annual Symposium Proceedings*, Nathanya, Israel, pp. 95–117.

Lua, W.F. (2007). Matching of Lowest Fare Seat Availability in Airline Revenue Management Systems. Master's Thesis, Massachusetts Institute of Technology, Cambridge, MA.

McCartney, S. (Feb 19, 2008). Unusual Route: Discount Airlines Woo Business Set. *Wall Street Journal*. Retrieved February 19, 2008, from the World Wide Web: http://online.wsj.com/article/SB120338017328875635.html?mod=dist_smartbrief.

McGill, J.I., G.J. van Ryzin. (1999). Revenue Management: Research Overviews and Prospects. *Transportation Science*, Volume 33, Issue 2, pp. 233-256.

Nason, S.D. (2006). Forecasting the future of airline revenue management. *Journal of Revenue and Pricing Management*, Volume 4, Issue 1, pp. 64-66.

Ozdaryal, K., Saranathan, B. (2004). "Revenue management in broken fare fence environment." *AGIFORS Reservations & Revenue Management Study Group Meeting*, Auckland, New Zealand.

Pickrell, D. (1991). The regulation and deregulation of US airlines. *Airline deregulation: international experiences.* Ed: Button, K. David Fulton Publishers, London, pp. 5-47.

Reyes, M. H. (2006). Hybrid Forecasting for Airline Revenue Management in Semi-Restricted Fare Structures. Master's Thesis, Massachusetts Institute of Technology, Cambridge, MA.

Rothstein, M. (1968). Stochastic Models for Airline Booking Policies. Ph.D. Thesis, Graduate School of Engineering and Science, New York University, New York, NY.

Rothstein, M. (1971). An airline overbooking model. *Transportation Science*, Volume 5, pp. 180-192.

Shumsky, R. (2006). The Southwest effect, airline alliances and revenue management. *Journal of Revenue and Pricing Management*, Volume 5, Issue 1, pp. 83-89.

Simpson, R. W. (1989). Using Network Flow Techniques to Find Shadow Prices for Market and Seat Inventory Control, Memorandum M89-1, MIT Flight Transportation Laboratory, Massachusetts Institute of Technology, Cambridge, MA.

Skwarek, D. K. (1996). Competitive Impacts of Yield Management Systems Components: Forecasting and Sell-up Models. Master's Thesis, Massachusetts Institute of Technology, Cambridge, MA.

Smith, B. C., C. W., Penn. (1988). Analysis of Alternative Origin-Destination Control Strategies, *AGIFORS Symposium Proceedings*, 28, 123-144.

Smith, B.C., J.F. Leimkuhler, R.M. Darrow. (1992). Yield Management at American Airlines. *Interfaces.* Volume 22, Issue 1, pp. 8-31.

Soo, Y.S.V. (2007). Fare Adjustment Strategies for Airline Revenue Management and Reservation Systems. Master's Thesis, Massachusetts Institute of Technology, Cambridge, MA.

Taylor, C. J. (1962). The Determination of Passenger Booking Levels. *2nd AGIFORS Annual Symposium Proceedings*, Fregene, Italy.

Thompson, H.R. (1961). Statistical Problems in Airline Reservations Control. *Operations Research Quarterly*. Volume 12, pp. 167–185.

Tretheway, M. W. (2004). Distortions of airline revenues: why the network airline business model is broken. *Journal of Air Transport Management*. Volume 10, Issue 1, pp. 3-14.

Usman, A. S. (2003). Demand forecasting accuracy in airline revenue management: analysis of practical issues with forecast error reduction. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

Vinod, B. (1995). Origin and Destination Yield Management. *The Handbook of Airline Economics*, D. Jenkins (ed.). The Aviation Weekly Group of the McGraw-Hill Companies, New York, NY, 459-468.

Weatherford, L. (1999). Forecast Aggregation and Disaggregation. *IATA Revenue Management Conference Proceedings*.

Weatherford, L.R. (1991). Perishable Asset Revenue Management in General Business Situations, Ph.D. thesis, Darden Graduate School of Business Administration, University of Virginia, Charlottesville, VA.

Weber, K., Thiel, R. (2004). "Optimisation issues in low cost revenue management." *AGIFORS Reservations & Revenue Management Study Group Meeting*, Auckland, New Zealand.

Wei, Y. J. (1997). Airline O-D Control using Network Displacement Concepts. Master's Thesis, Massachusetts Institute of Technology, Cambridge, MA.

Wickham, R. R. (1995). Evaluation of Forecasting Techniques for Short-Term Demand of Air Transportation. Master's Thesis, Massachusetts Institute of Technology, Cambridge, MA.

Williamson, E. L. (1988). Comparison of optimization techniques for origin-destination seat inventory control. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

Williamson, E. L. (1992). Airline Network Seat Inventory Control: Methodologies and Revenue Impacts. Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA.

Williamson, E. L. (1992). Airline Network Seat Inventory Control: Methodologies and Revenue Impacts. Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA.

Wollmer, R. D. (1992). An Airline Seat Management Model for a Single Leg Route when Lower Fare Classes Book First. *Operations Research.* Volume 40, Issue 1, pp. 26–37.

Zeni, R. H. (2001). Improved Forecast Accuracy in Revenue Management by Unconstraining Demand Estimates from Censored Data. Ph.D. Thesis. Rutgers, the State University of New Jersey, Newark, NJ.

Zickus, J. S. (1998). Forecasting for Airline Network Revenue Management: Revenue and Competitive Impacts. Master's Thesis, Massachusetts Institute of Technology, Cambridge, MA.